arXiv:1111.6082v3 [cs.LG] 27 Sep 2012

Trading Regret for Efficiency: Online Convex Optimization with Long Term Constraints

Mehrdad Mahdavi mahdavim@cse.msu.edu
Rong Jin rongjin@cse.msu.edu
Tianbao Yang yangtia1@msu.edu
Department of Computer Science and Engineering
Michigan State University
East Lansing, MI, 48824, USA

Editor:

Abstract

In this paper we propose efficient algorithms for solving constrained online convex optimization problems. Our motivation stems from the observation that most algorithms proposed for online convex optimization require a projection onto the convex set $\mathcal{K}$ from which the decisions are made. While the projection is straightforward for simple shapes (e.g., Euclidean ball), for arbitrary complex sets it is the main computational challenge and may be inefficient in practice. In this paper, we consider an alternative online convex optimization problem. Instead of requiring that decisions belong to $\mathcal{K}$ for all rounds, we only require that the constraints, which define the set $\mathcal{K}$, be satisfied in the long run. By turning the problem into an online convex-concave optimization problem, we propose an efficient algorithm which achieves $O(\sqrt{T})$ regret bound and $O(T^{3/4})$ bound on the violation of constraints. Then, we modify the algorithm in order to guarantee that the constraints are satisfied in the long run. This gain is achieved at the price of getting $O(T^{3/4})$ regret bound. Our second algorithm is based on the mirror prox method (Nemirovski, 2005) to solve variational inequalities which achieves $O(T^{2/3})$ bound for both regret and the violation of constraints when the domain $\mathcal{K}$ can be described by a finite number of linear constraints. Finally, we extend the results to the setting where we only have partial access to the convex set $\mathcal{K}$ and propose a multipoint bandit feedback algorithm with the same bounds in expectation as our first algorithm.

Keywords: online convex optimization, convex-concave optimization, bandit feedback, variational inequality

1. Introduction

Online convex optimization has recently emerged as a primitive framework for designing efficient algorithms for a wide variety of machine learning applications (Cesa-Bianchi and Lugosi, 2006). In general, an online convex optimization problem can be formulated as a repeated game between a learner and an adversary: at each iteration t, the learner first presents a solution $\mathbf{x}_t \in \mathcal{K}$, where $\mathcal{K} \subseteq \mathbb{R}^d$ is a convex domain representing the solution space; it then receives a convex function $f_t(\mathbf{x}) : \mathcal{K} \mapsto \mathbb{R}_+$ and suffers the loss $f_t(\mathbf{x}_t)$ for the submitted solution $\mathbf{x}_t$. The objective of the learner is to generate a sequence of solutions $\mathbf{x}_t \in \mathcal{K}, t = 1, 2, \cdots, T$ that minimizes the regret $\mathfrak{R}_T$ defined as

$$\mathfrak{R}_T = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - \min_{\mathbf{x}\in\mathcal{K}} \sum_{t=1}^{T} f_t(\mathbf{x}). \tag{1}$$

Regret measures the difference between the cumulative loss of the learner's strategy and the minimum possible loss had the sequence of loss functions been known in advance and the learner could choose the best fixed action in hindsight. When $\mathfrak{R}_T$ is sub-linear in the number of rounds, that is, $o(T)$, we call the solution Hannan consistent (Cesa-Bianchi and Lugosi, 2006), implying that the learner's average per-round loss approaches the average per-round loss of the best fixed action in hindsight. It is noticeable that the performance bound must hold for any sequence of loss functions, and in particular if the sequence is chosen adversarially.

Many successful algorithms have been developed over the past decade to minimize the regret in the online convex optimization. The problem was initiated in the remarkable work of Zinkevich (2003) which presents an algorithm based on gradient descent with projection that guarantees a regret of $O(\sqrt{T})$ when the set $\mathcal{K}$ is convex and the loss functions are Lipschitz continuous within the domain $\mathcal{K}$. In Hazan et al. (2007) and Shalev-Shwartz and Kakade (2008) algorithms with logarithmic regret bound were proposed for strongly convex loss functions. In particular, the algorithm in Hazan et al. (2007) is based on online Newton step and covers the general class of exp-concave loss functions. Notably, the simple gradient based algorithm also achieves an $O(\log T)$ regret bound for strongly convex loss functions with an appropriately chosen step size. Bartlett et al. (2007) generalizes the results in previous works to the setting where the algorithm can adapt to the curvature of the loss functions without any prior information. A modern view of these algorithms casts the problem as the task of following the regularized leader (Rakhlin, 2009). In Abernethy et al. (2009), using game-theoretic analysis, it has been shown that both $O(\sqrt{T})$ for Lipschitz continuous and $O(\log T)$ for strongly convex loss functions are tight in the minimax sense.

Examining the existing algorithms, most of the techniques usually require a projection step at each iteration in order to get back to the feasible region. For the performance of these online algorithms, the computational cost of the projection step is of crucial importance. To motivate the setting addressed in this paper, let us first examine a popular online learning algorithm for minimizing the regret $\mathfrak{R}_T$ based on the online gradient descent (OGD) method (Zinkevich, 2003). At each iteration t, after receiving the convex function $f_t(\mathbf{x})$, the learner computes the gradient $\nabla f_t(\mathbf{x}_t)$ and updates the solution $\mathbf{x}_t$ by solving the following optimization problem

$$\mathbf{x}_{t+1} = \Pi_{\mathcal{K}}\left(\mathbf{x}_t - \eta\nabla f_t(\mathbf{x}_t)\right) = \arg\min_{\mathbf{x}\in\mathcal{K}} \left\|\mathbf{x} - \mathbf{x}_t + \eta\nabla f_t(\mathbf{x}_t)\right\|^2, \tag{2}$$

where $\Pi_{\mathcal{K}}(\cdot)$ denotes the projection onto $\mathcal{K}$ and $\eta > 0$ is a predefined step size. Despite the simplicity of the OGD algorithm, the computational cost per iteration is crucial for its applicability. For general convex domains, solving the optimization problem in (2) is an offline convex optimization problem by itself and can be computationally expensive. For example, when one envisions a positive semidefinitive cone in applications such as distance metric learning and matrix completion, the full eigen-decomposition of a matrix is required to project the updated solutions back into the cone. Recently several efficient algorithms

have been developed for projection onto specific domains, for example, ℓ_1 ball (Duchi et al., 2008; Liu and Ye, 2009); however, when the domain $\mathcal{K}$ is complex, the projection step is a more involved task or computationally burdensome.

To tackle the computational challenge arising from the projection step, we consider an alternative online learning problem. Instead of requiring $\mathbf{x}_t \in \mathcal{K}$, we only require the constraints, which define the convex domain $\mathcal{K}$, to be satisfied in a long run. Then, the online learning problem becomes a task to find a sequence of solutions $\mathbf{x}_t, t \in [T]$ that minimizes the regret defined in (1), under the long term constraints, that is, $\sum_{t=1}^{T} \mathbf{x}_t/T \in \mathcal{K}$. We refer to this problem as **online learning with long term constraints**. In other words, instead of solving the projection problem in (2) on each round, we allow the learner to make decisions at some iterations which do not belong to the set $\mathcal{K}$, but the overall sequence of chosen decisions must obey the constraints at the end by a vanishing convergence rate.

From a different perspective, the proposed online optimization with long term constraints setup is reminiscent of regret minimization with side constraints or constrained regret minimization addressed in Mannor and Tsitsiklis (2006), motivated by applications in wireless communication. In regret minimization with side constraints, beyond minimizing regret, the learner has some side constraints that need to be satisfied on average for all rounds. Unlike our setting, in learning with side constraints, the set $\mathcal{K}$ is controlled by adversary and can vary arbitrarily from trial to trial. It has been shown that if the convex set is affected by both decisions and loss functions, the minimax optimal regret is generally unattainable online (Mannor et al., 2009).

One interesting application of the constrained regret minimization is multi-objective online classification where the learner aims at simultaneously optimizing more than one classification performance criteria. In the simple two objective online classification considered in Bernstein et al. (2010), the goal of the online classifier is to maximize the average true positive classification rate with an additional performance guarantee in terms of the false positive rate. Following the Neyman-Pearson risk, the intuitive approach to tackle this problem is to optimize one criterion (i.e., maximizing the true positive rate) subject to explicit constraint on the other objective (i.e., false positive rate) that needs to be satisfied on average for the sequence of decisions. The constrained regret matching (CRM) algorithm, proposed in Bernstein et al. (2010), efficiently solves this problem by relaxing the objective under mild assumptions on the single-stage constraint. The main idea of the CRM algorithm is to incorporate the penalty, that should be paid by the learner to satisfy the constraint, in the objective (i.e., true positive rate) by subtracting a positive constant at each decision step. It has been shown that the CRM algorithm *asymptotically* satisfies the average constraint (i.e., false positive rate) provided that the relaxation constant is above a certain threshold.

Finally, it is worth mentioning that the proposed setting can be used in certain classes of online learning such as online-to-batch conversion (Cesa-Bianchi et al., 2004), where it is sufficient to guarantee that constraints are satisfied in the long run. More specifically, under the assumption that received examples are i.i.d samples, the solution for batch learning is to average the solutions obtained over all the trials. As a result, if the long term constraint is satisfied, it is guaranteed that the average solution will belong to the domain $\mathcal{K}$.

In this paper, we describe and analyze a general framework for solving online convex optimization with long term constraints. We first show that a direct application of OGD

fails to achieve a sub-linear bound on the violation of constraints and an $O(\sqrt{T})$ bound on the regret. Then, by turning the problem into an online convex-concave optimization problem, we propose an efficient algorithm which is an adaption of OGD for online learning with long term constraints. The proposed algorithm achieves the same $O(\sqrt{T})$ regret bound as the general setting and $O(T^{3/4})$ bound for the violation of constraints. We show that by using a simple trick we can turn the proposed method into an algorithm which exactly satisfies the constraints in the long run by achieving $O(T^{3/4})$ regret bound. When the convex domain $\mathcal{K}$ can be described by a finite number of linear constraints, we propose an alternative algorithm based on the mirror prox method (Nemirovski, 2005), which achieves $O(T^{2/3})$ bound for both regret and the violation of constraints. Our framework also handles the cases when we do not have full access to the domain $\mathcal{K}$ except through a limited number of oracle evaluations. In the full-information version, the decision maker can observe the entire convex domain $\mathcal{K}$, whereas in a partial-information (a.k.a bandit setting) the decision maker may only observe the cost of the constraints defining the domain $\mathcal{K}$ at limited points. We show that we can generalize the proposed OGD based algorithm to this setting by only accessing the value oracle for domain $\mathcal{K}$ at two points, which achieves the same bounds in expectation as the case that has a full knowledge about the domain $\mathcal{K}$. In summary, the present work makes the following contributions:

- A general theorem that shows, in online setting, a simple penalty based method attains linear bound $O(T)$ for either the regret or the long term violation of the constraints and fails to achieve sub-linear bound for both regret and the long term violation of the constraints at the same time.
- A convex-concave formulation of online convex optimization with long term constraints, and an efficient algorithm based on OGD that attains a regret bound of $O(T^{1/2})$, and $O(T^{3/4})$ violation of the constraints.
- A modified OGD based algorithm for online convex optimization with long term constraints that has no constraint violation but $O(T^{3/4})$ regret bound.
- An algorithm for online convex optimization with long term constraints based on the mirror prox method that achieves $O(T^{2/3})$ regret and constraint violation.
- A multipoint bandit version of the basic algorithm with $O(T^{1/2})$ regret bound and $O(T^{3/4})$ violation of the constraints in *expectation* by accessing the value oracle for the convex set $\mathcal{K}$ at two points.

The remainder of the paper is structured as follows: In Section 3, we first examine a simple penalty based strategy and show that it fails to attain sub-linear bound for both regret and long term violation of the constraints. Then, we formulate regret minimization as an online convex-concave optimization problem and apply the OGD algorithm to solve it. Our first algorithm allows the constraints to be violated in a controlled way. It is then modified to have the constraints exactly satisfied in the long run. Section 4 presents our second algorithm which is an adaptation of the mirror prox method. Section 5 generalizes the online convex optimization with long term constraints problem to the setting where we only have a partial access to the convex domain $\mathcal{K}$. Section 6 concludes the work with a list of open questions.

2. Notation and Setting

Before proceeding, we define the notations used throughout the paper and state the assumptions made for the analysis of algorithms. Vectors are shown by lower case bold letters, such as $\mathbf{x} \in \mathbb{R}^d$. Matrices are indicated by upper case letters such as A and their pseudoinverse is represented by $A^\dagger$. We use $[m]$ as a shorthand for the set of integers $\{1, 2, \ldots, m\}$. Throughout the paper we denote by $\|\cdot\|$ and $\|\cdot\|_1$ the ℓ_2 (Euclidean) norm and ℓ_1-norm, respectively. We use $\mathbb{E}$ and $\mathbb{E}_t$ to denote the expectation and conditional expectation with respect to all randomness in early $t-1$ trials, respectively. To facilitate our analysis, we assume that the domain $\mathcal{K}$ can be written as an intersection of a finite number of convex constraints, that is,

$$\mathcal{K} = \{\mathbf{x} \in \mathbb{R}^d : g_i(\mathbf{x}) \leq 0, i \in [m]\},$$

where $g_i(\cdot), i \in [m]$, are Lipschitz continuous functions. Like many other works for online convex optimization such as Flaxman et al. (2005), we assume that $\mathcal{K}$ is a bounded domain, that is, there exist constants $R > 0$ and $r < 1$ such that $\mathcal{K} \subseteq R\mathbb{B}$ and $r\mathbb{B} \subseteq \mathcal{K}$ where $\mathbb{B}$ denotes the unit ℓ_2 ball centered at the origin. For the ease of notation, we use $\mathcal{B} = R\mathbb{B}$.

We focus on the problem of online convex optimization, in which the goal is to achieve a low regret with respect to a fixed decision on a sequence of loss functions. The difference between the setting considered here and the general online convex optimization is that, in our setting, instead of requiring $\mathbf{x}_t \in \mathcal{K}$, or equivalently $g_i(\mathbf{x}_t) \leq 0, i \in [m]$, for all $t \in [T]$, we only require the constraints to be satisfied in the long run, namely $\sum_{t=1}^T g_i(\mathbf{x}_t) \leq 0, i \in [m]$. Then, the problem becomes to find a sequence of solutions $\mathbf{x}_t, t \in [T]$ that minimizes the regret defined in (1), under the long term constraints $\sum_{t=1}^T g_i(\mathbf{x}_t) \leq 0, i \in [m]$. Formally, we would like to solve the following optimization problem online,

$$\min_{\mathbf{x}_1,\ldots,\mathbf{x}_T \in \mathcal{B}} \sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{\mathbf{x}\in\mathcal{K}} \sum_{t=1}^T f_t(\mathbf{x}) \;\text{ s.t. }\; \sum_{t=1}^T g_i(\mathbf{x}_t) \leq 0 \;,\; i \in [m]. \tag{3}$$

For simplicity, we will focus on a finite-horizon setting where the number of rounds T is known in advance. This condition can be relaxed under certain conditions, using standard techniques (see, e.g., Cesa-Bianchi and Lugosi, 2006). Note that in (3), (i) the solutions come from the ball $\mathcal{B} \supseteq \mathcal{K}$ instead of $\mathcal{K}$ and (ii) the constraint functions are fixed and are given in advance.

Like most online learning algorithms, we assume that both loss functions and the constraint functions are Lipschitz continuous, that is, there exists constants L_f and L_g such that $|f_t(\mathbf{x}) - f_t(\mathbf{x}')| \leq L_f\|\mathbf{x} - \mathbf{x}'\|, \quad |g_i(\mathbf{x}) - g_i(\mathbf{x}')| \leq L_g\|\mathbf{x} - \mathbf{x}'\|$ for any $\mathbf{x} \in \mathcal{B}$ and $\mathbf{x}' \in \mathcal{B}$, $i \in [m]$. For simplicity of analysis, we use $G = \max\{L_f, L_g\}$ and

$$\begin{aligned} F &= \max_{t\in[T]} \max_{\mathbf{x},\mathbf{x}'\in\mathcal{K}} f_t(\mathbf{x}) - f_t(\mathbf{x}') \leq 2L_f R, \\ D &= \max_{i\in[m]} \max_{\mathbf{x}\in\mathcal{B}} g_i(\mathbf{x}) \leq L_g R. \end{aligned}$$

Finally, we define the notion of a Bregman divergence. Let $\phi(\cdot)$ be a strictly convex function defined on a convex set $\mathcal{K}$. The Bregman divergence between $\mathbf{x}$ and $\mathbf{x}'$ is defined as $\mathsf{B}_\phi(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x}) - \phi(\mathbf{x}') - (\mathbf{x} - \mathbf{x}')^\top \nabla\phi(\mathbf{x}')$ which measures how much the function $\phi(\cdot)$ deviates at $\mathbf{x}$ from it's linear approximation at $\mathbf{x}'$.

3. Online Convex Optimization with Long Term Constraints

In this section we present and analyze our gradient descent based algorithms for online convex optimization problem with long term constraints. We first describe an algorithm which is allowed to violate the constraints and then, by applying a simple trick, we propose a variant of the first algorithm which exactly satisfies the constraints in the long run.

Before we state our formulation and algorithms, let us review a few alternative techniques that do not need explicit projection. A straightforward approach is to introduce an appropriate self-concordant barrier function for the given convex set $\mathcal{K}$ and add it to the objective function such that the barrier diverges at the boundary of the set. Then we can interpret the resulting optimization problem, on the modified objective functions, as an unconstrained minimization problem that can be solved without projection steps. Following the analysis in Abernethy et al. (2012), with an appropriately designed procedure for updating solutions, we could guarantee a regret bound of $O(\sqrt{T})$ without the violation of constraints. A similar idea is used in Abernethy et al. (2008) for online bandit learning and in Narayanan and Rakhlin (2010) for a random walk approach for regret minimization which, in fact, translates the issue of projection into the difficulty of sampling. Even for linear Lipschitz cost functions, the random walk approach requires sampling from a Gaussian distribution with covariance given by the Hessian of the self-concordant barrier of the convex set $\mathcal{K}$ that has the same time complexity as inverting a matrix. The main limitation with these approaches is that they require computing the Hessian matrix of the objective function in order to guarantee that the updated solution stays within the given domain $\mathcal{K}$. This limitation makes it computationally unattractive when dealing with high dimensional data. In addition, except for well known cases, it is often unclear how to efficiently construct a self-concordant barrier function for a general convex domain.

An alternative approach for online convex optimization with long term constraints is to introduce a penalty term in the loss function that penalizes the violation of constraints. More specifically, we can define a new loss function $\hat{f}_t(\cdot)$ as

$$\hat{f}_t(\mathbf{x}) = f_t(\mathbf{x}) + \delta \sum_{i=1}^{m} [g_i(\mathbf{x})]_+, \tag{4}$$

where $[z]_+ = \max(0, 1 - z)$ and $\delta > 0$ is a fixed positive constant used to penalize the violation of constraints. We then run the standard OGD algorithm to minimize the modified loss function $\hat{f}_t(\cdot)$. The following theorem shows that this simple strategy fails to achieve sub-linear bound for both regret and the long term violation of constraints at the same time.

Theorem 1 *Given $\delta > 0$, there always exists a sequence of loss functions $\{f_t(\mathbf{x})\}_{t=1}^T$ and a constraint function $g(\mathbf{x})$ such that either $\sum_{t=1}^T f_t(\mathbf{x}_t) - \min_{g(\mathbf{x})\leq 0} \sum_{t=1}^T f_t(\mathbf{x}) = O(T)$ or $\sum_{t=1}^T [g(\mathbf{x}_t)]_+ = O(T)$ holds, where $\{\mathbf{x}_t\}_{t=1}^T$ is the sequence of solutions generated by the OGD algorithm that minimizes the modified loss functions given in (4).*

We defer the proof to Appendix A along with a simple analysis of the OGD when applied to the modified functions in (4). The analysis shows that in order to obtain $O(\sqrt{T})$ regret bound, linear bound on the long term violation of the constraints is unavoidable. The main

reason for the failure of using modified loss function in (4) is that the weight constant δ is fixed and independent from the sequence of solutions obtained so far. In the next subsection, we present an online convex-concave formulation for online convex optimization with long term constraints, which explicitly addresses the limitation of (4) by automatically adjusting the weight constant based on the violation of the solutions obtained so far.

As mentioned before, our general strategy is to turn online convex optimization with long term constraints into a convex-concave optimization problem. Instead of generating a sequence of solutions that satisfies the long term constraints, we first consider an online optimization strategy that allows the violation of constraints on some rounds in a controlled way. We then modify the online optimization strategy to obtain a sequence of solutions that obeys the long term constraints. Although the online convex optimization with long term constraints is clearly easier than the standard online convex optimization problem, it is straightforward to see that optimal regret bound for online optimization with long term constraints should be on the order of $O(\sqrt{T})$, no better than the standard online convex optimization problem.

3.1 An Efficient Algorithm with $O(\sqrt{T})$ Regret Bound and $O(T^{3/4})$ Bound on the Violation of Constraints

The intuition behind our approach stems from the observation that the constrained optimization problem $\min_{\mathbf{x}\in\mathcal{K}} \sum_{t=1}^{T} f_t(\mathbf{x})$ is equivalent to the following convex-concave optimization problem

$$\min_{\mathbf{x}\in\mathcal{B}} \max_{\lambda\in\mathbb{R}_+^m} \sum_{t=1}^{T} f_t(\mathbf{x}) + \sum_{i=1}^{m} \lambda_i g_i(\mathbf{x}), \tag{5}$$

where $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_m)^\top$ is the vector of Lagrangian multipliers associated with the constraints $g_i(\cdot), i = 1, \ldots, m$ and belongs to the nonnegative orthant $\mathbb{R}_+^m$. To solve the online convex-concave optimization problem, we extend the gradient based approach for variational inequality (Nemirovski, 1994) to (5). To this end, we consider the following regularized convex-concave function as

$$\mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}) = f_t(\mathbf{x}) + \sum_{i=1}^{m} \left\{ \lambda_i g_i(\mathbf{x}) - \frac{\delta\eta}{2}\lambda_i^2 \right\}, \tag{6}$$

where $\delta > 0$ is a constant whose value will be decided by the analysis. Note that in (6), we introduce a regularizer $\delta\eta\lambda_i^2/2$ to prevent λ_i from being too large. This is because, when λ_i is large, we may encounter a large gradient for $\mathbf{x}$ because of $\nabla_{\mathbf{x}}\mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}) \propto \sum_{i=1}^{m} \lambda_i \nabla g_i(\mathbf{x})$, leading to unstable solutions and a poor regret bound. Although we can achieve the same goal by restricting λ_i to a bounded domain, using the quadratic regularizer makes it convenient for our analysis.

Algorithm 1 shows the detailed steps of the proposed algorithm. Unlike standard online convex optimization algorithms that only update $\mathbf{x}$, Algorithm 1 updates both $\mathbf{x}$ and $\boldsymbol{\lambda}$. In addition, unlike the modified loss function in (4) where the weights for constraints $\{g_i(\mathbf{x}) \leq 0\}_{i=1}^m$ are fixed, Algorithm 1 automatically adjusts the weights $\{\lambda_i\}_{i=1}^m$ based on $\{g_i(\mathbf{x})\}_{i=1}^m$, the violation of constraints, as the game proceeds. It is this property that allows Algorithm 1 to achieve sub-linear bound for both regret and the violation of constraints.

Algorithm 1 Gradient based Online Convex Optimization with Long Term Constraints

Input: constraints $g_i(\mathbf{x}) \leq 0, i \in [m]$, step size η, and constant $\delta > 0$
Initialization: $\mathbf{x}_1 = \mathbf{0}$ and $\boldsymbol{\lambda}_1 = \mathbf{0}$
for $t = 1, 2, \ldots, T$ **do**
  Submit solution $\mathbf{x}_t$
  Receive the convex function $f_t(\mathbf{x})$ and experience loss $f_t(\mathbf{x}_t)$
  Compute $\nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t) = \nabla f_t(\mathbf{x}_t) + \sum_{i=1}^m \lambda_t^i \nabla g_i(\mathbf{x}_t)$ and $\nabla_{\lambda_i} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t) = g_i(\mathbf{x}_t) - \eta\delta\lambda_t^i$
  Update $\mathbf{x}_t$ and $\boldsymbol{\lambda}_t$ by

$$\mathbf{x}_{t+1} = \Pi_{\mathcal{B}} \left(\mathbf{x}_t - \eta \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t) \right)$$
$$\boldsymbol{\lambda}_{t+1} = \Pi_{[0,+\infty)^m} (\boldsymbol{\lambda}_t + \eta \nabla_{\boldsymbol{\lambda}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t))$$

end for

To analyze Algorithm 1, we first state the following lemma, the key to the main theorem on the regret bound and the violation of constraints.

Lemma 2 *Let $\mathcal{L}_t(\cdot, \cdot)$ be the function defined in (6) which is convex in its first argument and concave in its second argument. Then for any $(\mathbf{x}, \boldsymbol{\lambda}) \in \mathcal{B} \times \mathbb{R}_+^m$ we have*

$$\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}) - \mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}_t) \leq \frac{1}{2\eta}(\|\mathbf{x} - \mathbf{x}_t\|^2 + \|\boldsymbol{\lambda} - \boldsymbol{\lambda}_t\|^2 - \|\mathbf{x} - \mathbf{x}_{t+1}\|^2 - \|\boldsymbol{\lambda} - \boldsymbol{\lambda}_{t+1}\|^2) + \frac{\eta}{2}(\|\nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2 + \|\nabla_{\boldsymbol{\lambda}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2).$$

Proof Following the analysis of Zinkevich (2003), convexity of $\mathcal{L}_t(\cdot, \boldsymbol{\lambda})$ implies that

$$\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t) - \mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}_t) \leq (\mathbf{x}_t - \mathbf{x})^\top \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t) \tag{7}$$

and by concavity of $\mathcal{L}_t(\mathbf{x}, \cdot)$ we have

$$\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}) - \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t) \leq (\boldsymbol{\lambda} - \boldsymbol{\lambda}_t)^\top \nabla_{\boldsymbol{\lambda}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t). \tag{8}$$

Combining the inequalities (7) and (8) results in

$$\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}) - \mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}_t) \leq (\mathbf{x}_t - \mathbf{x})^\top \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t) - (\boldsymbol{\lambda} - \boldsymbol{\lambda}_t)^\top \nabla_{\boldsymbol{\lambda}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t). \tag{9}$$

Using the update rule for $\mathbf{x}_{t+1}$ in terms of $\mathbf{x}_t$ and expanding, we get

$$\|\mathbf{x} - \mathbf{x}_{t+1}\|^2 \leq \|\mathbf{x} - \mathbf{x}_t\|^2 - 2\eta(\mathbf{x}_t - \mathbf{x})^\top \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t) + \eta^2 \|\nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2, \tag{10}$$

where the first inequality follows from the nonexpansive property of the projection operation. Expanding the inequality for $\|\boldsymbol{\lambda} - \boldsymbol{\lambda}_{t+1}\|^2$ in terms of $\boldsymbol{\lambda}_t$ and plugging back into the (9) with (10) establishes the desired inequality. ■

Proposition 3 *Let $\mathbf{x}_t$ and $\boldsymbol{\lambda}_t, t \in [T]$ be the sequence of solutions obtained by Algorithm 1. Then for any $\mathbf{x} \in \mathcal{B}$ and $\boldsymbol{\lambda} \in \mathbb{R}_+^m$, we have*

$$\sum_{t=1}^{T} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}) - \mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}_t) \tag{11}$$
$$\leq \frac{R^2 + \|\boldsymbol{\lambda}\|^2}{2\eta} + \frac{\eta T}{2}\left((m+1)G^2 + 2mD^2\right) + \frac{\eta}{2}\left((m+1)G^2 + 2m\delta^2\eta^2\right)\sum_{t=1}^{T}\|\boldsymbol{\lambda}_t\|^2.$$

Proof We first bound the gradient terms in the right hand side of Lemma 2. Using the inequality $(a_1 + a_2 + \ldots, a_n)^2 \leq n(a_1^2 + a_2^2 + \ldots + a_n^2)$, we have $\|\nabla_{\mathbf{x}}\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2 \leq (m+1)G^2\left(1 + \|\boldsymbol{\lambda}_t\|^2\right)$ and $\|\nabla_{\boldsymbol{\lambda}}\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2 \leq 2m(D^2 + \delta^2\eta^2\|\boldsymbol{\lambda}_t\|^2)$. In Lemma 2, by adding the inequalities of all iterations, and using the fact $\|\mathbf{x}\| \leq R$ we complete the proof. ■

The following theorem bounds the regret and the violation of the constraints in the long run for Algorithm 1.

Theorem 4 *Define $a = R\sqrt{(m+1)G^2 + 2mD^2}$. Set $\eta = R^2/[a\sqrt{T}]$. Assume T is large enough such that $2\sqrt{2}\eta(m+1) \leq 1$. Choose δ such that $\delta \geq (m+1)G^2 + 2m\delta^2\eta^2$. Let $\mathbf{x}_t, t \in [T]$ be the sequence of solutions obtained by Algorithm 1. Then for the optimal solution $\mathbf{x}_* = \min_{\mathbf{x}\in\mathcal{K}}\sum_{t=1}^{T} f_t(\mathbf{x})$ we have*

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) \leq a\sqrt{T} = O(T^{1/2}), \text{ and}$$
$$\sum_{t=1}^{T} g_i(\mathbf{x}_t) \leq \sqrt{2\left(FT + a\sqrt{T}\right)\sqrt{T}\left(\frac{\delta R^2}{a} + \frac{ma}{R^2}\right)} = O(T^{3/4}).$$

Proof We begin by expanding (11) using (6) and rearranging the terms to get

$$\sum_{t=1}^{T}[f_t(\mathbf{x}_t) - f_t(\mathbf{x})] + \sum_{i=1}^{m}\left\{\lambda_i\sum_{t=1}^{T} g_i(\mathbf{x}_t) - \sum_{t=1}^{T}\lambda_t^i g_i(\mathbf{x})\right\} - \frac{\delta\eta T}{2}\|\boldsymbol{\lambda}\|^2$$
$$\leq -\frac{\delta\eta}{2}\sum_{t=1}^{T}\|\boldsymbol{\lambda}_t\|^2 + \frac{R^2 + \|\boldsymbol{\lambda}\|^2}{2\eta} + \frac{\eta T}{2}\left((m+1)G^2 + 2mD^2\right)$$
$$+ \frac{\eta}{2}\left((m+1)G^2 + 2m\delta^2\eta^2\right)\sum_{t=1}^{T}\|\boldsymbol{\lambda}_t\|^2.$$

Since $\delta \geq (m+1)G^2 + 2m\delta^2\eta^2$, we can drop the $\|\boldsymbol{\lambda}_t\|^2$ terms from both sides of the above inequality and obtain

$$\sum_{t=1}^{T}[f_t(\mathbf{x}_t) - f_t(\mathbf{x})] + \sum_{i=1}^{m}\left\{\lambda_i\sum_{t=1}^{T} g_i(\mathbf{x}_t) - \left(\frac{\delta\eta T}{2} + \frac{m}{2\eta}\right)\lambda_i^2\right\}$$
$$\leq \sum_{i=1}^{m}\sum_{t=1}^{T}\lambda_t^i g_i(\mathbf{x}) + \frac{R^2}{2\eta} + \frac{\eta T}{2}\left((m+1)G^2 + 2mD^2\right).$$

The left hand side of above inequality consists of two terms. The first term basically measures the difference between the cumulative loss of the Algorithm 1 and the optimal solution and the second term includes the constraint functions with corresponding Lagrangian multipliers which will be used to bound the long term violation of the constraints. By taking maximization for $\boldsymbol{\lambda}$ over the range $(0, +\infty)$, we get

$$\sum_{t=1}^{T} [f_t(\mathbf{x}_t) - f_t(\mathbf{x})] + \sum_{i=1}^{m} \left\{ \frac{\left[\sum_{t=1}^{T} g_i(\mathbf{x}_t)\right]_+^2}{2(\delta\eta T + m/\eta)} - \sum_{t=1}^{T} \lambda_t^i g_i(\mathbf{x}) \right\} \leq \frac{R^2}{2\eta} + \frac{\eta T}{2}\left((m+1)G^2 + 2mD^2)\right).$$

Since $\mathbf{x}_* \in \mathcal{K}$, we have $g_i(\mathbf{x}_*) \leq 0, i \in [m]$, and the resulting inequality becomes

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) + \sum_{i=1}^{m} \frac{\left[\sum_{t=1}^{T} g_i(\mathbf{x}_t)\right]_+^2}{2(\delta\eta T + m/\eta)} \leq \frac{R^2}{2\eta} + \frac{\eta T}{2}\left((m+1)G^2 + 2mD^2)\right).$$

The statement of the first part of the theorem follows by using the expression for η. The second part is proved by substituting the regret bound by its lower bound as $\sum_{t=1}^{T} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) \geq -FT$. ■

Remark 5 *We observe that the introduction of quadratic regularizer $\delta\eta\|\boldsymbol{\lambda}\|^2/2$ allows us to turn the expression $\lambda_i \sum_{t=1}^{T} g_i(\mathbf{x}_t)$ into $\left[\sum_{t=1}^{T} g_i(\mathbf{x}_t)\right]_+^2$, leading to the bound for the violation of the constraints. In addition, the quadratic regularizer defined in terms of $\boldsymbol{\lambda}$ allows us to work with unbounded $\boldsymbol{\lambda}$ because it cancels the contribution of the $\|\boldsymbol{\lambda}_t\|$ terms from the loss function and the bound on the gradients $\|\nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda})\|$. Note that the constraint for δ mentioned in Theorem 4 is equivalent to*

$$\frac{2}{1/(m+1) + \sqrt{(m+1)^{-2} - 8G^2\eta^2}} \leq \delta \leq \frac{1/(m+1) + \sqrt{(m+1)^{-2} - 8G^2\eta^2}}{4\eta^2}, \qquad (12)$$

from which, when T is large enough (i.e., η is small enough), we can simply set $\delta = 2(m+1)G^2$ that will obey the constraint in (12).

By investigating Lemma 2, it turns out that the boundedness of the gradients is essential to obtain bounds for Algorithm 1 in Theorem 4. Although, at each iteration, $\boldsymbol{\lambda}_t$ is projected onto the $\mathbb{R}_+^m$, since $\mathcal{K}$ is a compact set and functions $f_t(\mathbf{x})$ and $g_i(\mathbf{x}), i \in [m]$ are convex, the boundedness of the functions implies that the gradients are bounded (Bertsekas et al., 2003, Proposition 4.2.3).

3.2 An Efficient Algorithm with $O(T^{3/4})$ Regret Bound and without Violation of Constraints

In this subsection we generalize Algorithm 1 such that the constrained are satisfied in a long run. To create a sequence of solutions $\{\mathbf{x}_t, t \in [T]\}$ that satisfies the long term

constraints $\sum_{t=1}^{T} g_i(\mathbf{x}_t) \leq 0, i \in [m]$, we make two modifications to Algorithm 1. First, instead of handling all of the m constraints, we consider a single constraint defined as $g(\mathbf{x}) = \max_{i\in[m]} g_i(\mathbf{x})$. Apparently, by achieving zero violation for the constraint $g(\mathbf{x}) \leq 0$, it is guaranteed that all of the constraints $g_i(\cdot), i \in [m]$ are also satisfied in the long term. Furthermore, we change Algorithm 1 by modifying the definition of $\mathcal{L}_t(\cdot,\cdot)$ as

$$\mathcal{L}_t(\mathbf{x}, \lambda) = f_t(\mathbf{x}) + \lambda(g(\mathbf{x}) + \gamma) - \frac{\eta\delta}{2}\lambda^2, \tag{13}$$

where $\gamma > 0$ will be decided later. This modification is equivalent to considering the constraint $g(\mathbf{x}) \leq -\gamma$, a tighter constraint than $g(\mathbf{x}) \leq 0$. The main idea behind this modification is that by using a tighter constraint in our algorithm, the resulting sequence of solutions will satisfy the long term constraint $\sum_{t=1}^{T} g(\mathbf{x}_t) \leq 0$, even though the tighter constraint is violated in many trials.

Before proceeding, we state a fact about the Lipschitz continuity of the function $g(\mathbf{x})$ in the following proposition.

Proposition 6 *Assume that functions $g_i(\cdot), i \in [m]$ are Lipschitz continuous with constant G. Then, function $g(\mathbf{x}) = \max_{i\in[m]} g_i(\mathbf{x})$ is Lipschitz continuous with constant G, that is,*

$$|g(\mathbf{x}) - g(\mathbf{x}')| \leq G\|\mathbf{x} - \mathbf{x}'\| \;\; \textit{for any } \mathbf{x} \in \mathcal{B} \textit{ and } \mathbf{x}' \in \mathcal{B}.$$

Proof First, we rewrite $g(\mathbf{x}) = \max_{i\in[m]} g_i(\mathbf{x})$ as $g(\mathbf{x}) = \max_{\alpha\in\Delta_m} \sum_{i=1}^{m} \alpha_i g_i(\mathbf{x})$ where Δ_m is the m-simplex, that is, $\Delta_m = \{\alpha \in \mathbb{R}_+^m; \sum_{i=1}^{m} \alpha_i = 1\}$. Then, we have

$$\begin{aligned}
|g(\mathbf{x}) - g(\mathbf{x}')| &= \left| \max_{\alpha\in\Delta_m} \sum_{i=1}^{m} \alpha_i g_i(\mathbf{x}) - \max_{\alpha\in\Delta_m} \sum_{i=1}^{m} \alpha_i g_i(\mathbf{x}') \right| \\
&\leq \max_{\alpha\in\Delta_m} \left| \sum_{i=1}^{m} \alpha_i g_i(\mathbf{x}) - \sum_{i=1}^{m} \alpha_i g_i(\mathbf{x}') \right| \\
&\leq \max_{\alpha\in\Delta_m} \sum_{i=1}^{m} \alpha_i \left| g_i(\mathbf{x}) - g_i(\mathbf{x}') \right| \leq G\|\mathbf{x} - \mathbf{x}'\|,
\end{aligned}$$

where the last inequality follows from the Lipschitz continuity of $g_i(\mathbf{x}), i \in [m]$. ■

To obtain a zero bound on the violation of constraints in the long run, we make the following assumption about the constraint function $g(\mathbf{x})$.

Assumption 1 *Let $\mathcal{K}' \subseteq \mathcal{K}$ be the convex set defined as $\mathcal{K}' = \{\mathbf{x} \in \mathbb{R}^d : g(\mathbf{x}) + \gamma \leq 0\}$ where $\gamma \geq 0$. We assume that the norm of the gradient of the constraint function $g(\mathbf{x})$ is lower bounded at the boundary of $\mathcal{K}'$, that is,*

$$\min_{g(\mathbf{x})+\gamma=0} \|\nabla g(\mathbf{x})\| \geq \sigma.$$

A direct consequence of **Assumption 1** is that by reducing the domain $\mathcal{K}$ to $\mathcal{K}'$, the optimal value of the constrained optimization problem $\min_{\mathbf{x}\in\mathcal{K}} f(\mathbf{x})$ does not change much, as revealed by the following theorem.

Theorem 7 *Let $\mathbf{x}_*$ and $\mathbf{x}_\gamma$ be the optimal solutions to the constrained optimization problems defined as $\min_{g(\mathbf{x})\leq 0} f(\mathbf{x})$ and $\min_{g(\mathbf{x})\leq -\gamma} f(\mathbf{x})$, respectively, where $f(\mathbf{x}) = \sum_{t=1}^T f_t(\mathbf{x})$ and $\gamma \geq 0$. We have*

$$|f(\mathbf{x}_*) - f(\mathbf{x}_\gamma)| \leq \frac{G}{\sigma}\gamma T.$$

Proof We note that the optimization problem $\min_{g(\mathbf{x})\leq -\gamma} f(\mathbf{x}) = \min_{g(\mathbf{x})\leq -\gamma} \sum_{t=1}^T f_t(\mathbf{x})$, can also be written in the minimax form as

$$f(\mathbf{x}_\gamma) = \min_{\mathbf{x}\in\mathcal{B}} \max_{\lambda\in\mathbb{R}_+} \sum_{t=1}^T f_t(\mathbf{x}) + \lambda(g(\mathbf{x}) + \gamma), \tag{14}$$

where we use the fact that $\mathcal{K}' \subseteq \mathcal{K} \subseteq \mathcal{B}$. We denote by $\mathbf{x}_\gamma$ and λ_γ the optimal solutions to (14). We have

$$\begin{aligned} f(\mathbf{x}_\gamma) &= \min_{\mathbf{x}\in\mathcal{B}} \max_{\lambda\in\mathbb{R}_+} \sum_{t=1}^T f_t(\mathbf{x}) + \lambda(g(\mathbf{x}) + \gamma) \\ &= \min_{\mathbf{x}\in\mathcal{B}} \sum_{t=1}^T f_t(\mathbf{x}) + \lambda_\gamma(g(\mathbf{x}) + \gamma) \\ &\leq \sum_{t=1}^T f_t(\mathbf{x}_*) + \lambda_\gamma(g(\mathbf{x}_*) + \gamma) \leq \sum_{t=1}^T f_t(\mathbf{x}_*) + \lambda_\gamma\gamma, \end{aligned}$$

where the second equality follows the definition of the $\mathbf{x}_\gamma$ and the last inequality is due to the optimality of $\mathbf{x}_*$, that is, $g(\mathbf{x}_*) \leq 0$.
To bound $|f(\mathbf{x}_\gamma) - f(\mathbf{x}_*)|$, we need to bound λ_γ. Since $\mathbf{x}_\gamma$ is the minimizer of (14), from the optimality condition we have

$$-\sum_{t=1}^T \nabla f_t(\mathbf{x}_\gamma) = \lambda_\gamma \nabla g(\mathbf{x}_\gamma). \tag{15}$$

By setting $\mathbf{v} = -\sum_{t=1}^T \nabla f_t(\mathbf{x}_\gamma)$, we can simplify (15) as $\lambda_\gamma \nabla g(\mathbf{x}_\gamma) = \mathbf{v}$. From the KKT optimality condition (Boyd and Vandenberghe, 2004), if $g(\mathbf{x}_\gamma)+\gamma < 0$ then we have $\lambda_\gamma = 0$; otherwise according to Assumption 1 we can bound λ_γ by

$$\lambda_\gamma \leq \frac{\|\mathbf{v}\|}{\|\nabla g(\mathbf{x}_\gamma)\|} \leq \frac{GT}{\sigma}.$$

We complete the proof by applying the fact $f(\mathbf{x}_*) \leq f(\mathbf{x}_\gamma) \leq f(\mathbf{x}_*) + \lambda_\gamma\gamma$. ■

As indicated by Theorem 7, when γ is small, we expect the difference between two optimal values $f(\mathbf{x}_*)$ and $f(\mathbf{x}_\gamma)$ to be small. Using the result from Theorem 7, in the following theorem, we show that by running Algorithm 1 on the modified convex-concave functions defined in (13), we are able to obtain an $O(T^{3/4})$ regret bound and zero bound on the violation of constraints in the long run.

Theorem 8 *Set* $a = 2R/\sqrt{2G^2 + 3(D^2 + b^2)}$, $\eta = R^2/[a\sqrt{T}]$, *and* $\delta = 4G^2$. *Let* $\mathbf{x}_t, t \in [T]$ *be the sequence of solutions obtained by Algorithm 1 with functions defined in (13) with* $\gamma = bT^{-1/4}$ *and* $b = 2\sqrt{F(\delta R^2 a^{-1} + aR^{-2})}$. *Let* $\mathbf{x}_*$ *be the optimal solution to* $\min_{\mathbf{x}\in\mathcal{K}} \sum_{t=1}^T f_t(\mathbf{x})$. *With sufficiently large* T, *that is,* $FT \geq a\sqrt{T}$, *and under Assumption 1, we have* $\mathbf{x}_t, t \in [T]$ *satisfy the global constraint* $\sum_{t=1}^T g(\mathbf{x}_t) \leq 0$ *and the regret* $\mathfrak{R}_T$ *is bounded by*

$$\mathfrak{R}_T = \sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) \leq a\sqrt{T} + \frac{b}{\sigma} GT^{3/4} = O(T^{3/4}).$$

Proof Let $\mathbf{x}_\gamma$ be the optimal solution to $\min_{g(\mathbf{x})\leq -\gamma} \sum_{t=1}^T f_t(\mathbf{x})$. Similar to the proof of Theorem 4 when applied to functions in (13) we have

$$\begin{aligned}
&\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}) + \lambda \sum_{t=1}^T (g(\mathbf{x}_t) + \gamma) - \left(\sum_{t=1}^T \lambda_t\right)(g(\mathbf{x}) + \gamma) - \frac{\delta\eta T}{2}\lambda^2 \\
&\leq \quad -\frac{\delta\eta}{2}\sum_{t=1}^T \lambda_t^2 + \frac{R^2 + \lambda^2}{2\eta} + \frac{\eta T}{2}\left(2G^2 + 3(D^2 + \gamma^2)\right) + \frac{\eta}{2}\left(2G^2 + 3\delta^2\eta^2\right)\sum_{t=1}^T \lambda_t^2.
\end{aligned}$$

By setting $\delta \geq 2G^2 + 3\delta^2\eta^2$ which is satisfied by $\delta = 4G^2$, we cancel the terms including λ_t from the right hand side of above inequality. By maximizing for λ over the range $(0, +\infty)$ and noting that $\gamma \leq b$, for the optimal solution $\mathbf{x}_\gamma$, we have

$$\sum_{t=1}^T [f_t(\mathbf{x}_t) - f_t(\mathbf{x}_\gamma)] + \frac{\left[\sum_{t=1}^T g(\mathbf{x}_t) + \gamma T\right]_+^2}{2(\delta\eta T + 1/\eta)} \leq \frac{R^2}{2\eta} + \frac{\eta T}{2}\left(2G^2 + 3(D^2 + b^2)\right),$$

which, by optimizing for η and applying the lower bound for the regret as $\sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_\gamma) \geq -FT$, yields the following inequalities

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_\gamma) \leq a\sqrt{T} \tag{16}$$

and

$$\sum_{t=1}^T g(\mathbf{x}_t) \leq \sqrt{2\left(FT + a\sqrt{T}\right)\sqrt{T}\left(\frac{\delta R^2}{a} + \frac{a}{R^2}\right)} - \gamma T, \tag{17}$$

for the regret and the violation of the constraint, respectively. Combining (16) with the result of Theorem 7 results in $\sum_{t=1}^T f_t(\mathbf{x}_\gamma) \leq \sum_{t=1}^T f_t(\mathbf{x}_*) + a\sqrt{T} + (G/\sigma)\gamma T$. By choosing $\gamma = bT^{-1/4}$ we attain the desired regret bound as

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) \leq a\sqrt{T} + \frac{bG}{\sigma} T^{3/4} = O(T^{3/4}).$$

To obtain the bound on the violation of constraints, we note that in (17), when T is sufficiently large, that is, $FT \geq a\sqrt{T}$, we have $\sum_{t=1}^T g(\mathbf{x}_t) \leq 2\sqrt{F(\delta R^2 a^{-1} + aR^{-2})}T^{3/4} - bT^{3/4}$.

Choosing $b = 2\sqrt{F(\delta R^2 a^{-1} + aR^{-2})}T^{3/4}$ guarantees the zero bound on the violation of constraints as claimed. ■

4. A Mirror Prox Based Approach

The bound for the violation of constraints for Algorithm 1 is unsatisfactory since it is significantly worse than $O(\sqrt{T})$. In this section, we pursue a different approach that is based on the mirror prox method in Nemirovski (2005) to improve the bound for the violation of constraints. The basic idea is that solving (5) can be reduced to the problem of approximating a saddle point $(\mathbf{x}, \boldsymbol{\lambda}) \in \mathcal{B} \times [0, \infty)^m$ by solving the associated variational inequality.

We first define an auxiliary function $\mathcal{F}(\mathbf{x}, \boldsymbol{\lambda})$ as

$$\mathcal{F}(\mathbf{x}, \boldsymbol{\lambda}) = \sum_{i=1}^{m} \left\{ \lambda_i g_i(\mathbf{x}) - \frac{\delta\eta}{2}\lambda_i^2 \right\}.$$

In order to successfully apply the mirror prox method, we follow the fact that any convex domain can be written as an intersection of linear constraints, and make the following assumption:

Assumption 2 *We assume that $g_i(\mathbf{x}), i \in [m]$ are linear, that is, $\mathcal{K} = \{\mathbf{x} \in \mathbb{R}^d : g_i(\mathbf{x}) = \mathbf{x}^\top \mathbf{a}_i - b_i \leq 0, i \in [m]\}$ where $\mathbf{a}_i \in \mathbb{R}^d$ is a normalized vector with $\|\mathbf{a}_i\| = 1$ and $b_i \in \mathbb{R}$.*

The following proposition shows that under Assumptions 2, the function $\mathcal{F}(\mathbf{x}, \boldsymbol{\lambda})$ has Lipschitz continuous gradient, a basis for the application of the mirror prox method.

Proposition 9 *Under Assumption 2, $\mathcal{F}(\mathbf{x}, \boldsymbol{\lambda})$ has Lipschitz continuous gradient, that is,*

$$\left\|\nabla_{\mathbf{x}}\mathcal{F}(\mathbf{x}, \boldsymbol{\lambda}) - \nabla_{\mathbf{x}'}\mathcal{F}(\mathbf{x}', \boldsymbol{\lambda}')\right\|^2 + \left\|\nabla_{\boldsymbol{\lambda}}\mathcal{F}(\mathbf{x}, \boldsymbol{\lambda}) - \nabla_{\boldsymbol{\lambda}'}\mathcal{F}(\mathbf{x}', \boldsymbol{\lambda}')\right\|^2 \leq 2(m + \delta^2\eta^2)(\|\mathbf{x} - \mathbf{x}'\|^2 + \|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\|^2).$$

Proof

$$\begin{aligned}
& \left\|\nabla_{\mathbf{x}}\mathcal{F}(\mathbf{x}, \boldsymbol{\lambda}) - \nabla_{\mathbf{x}'}\mathcal{F}(\mathbf{x}', \boldsymbol{\lambda}')\right\|^2 + \left\|\nabla_{\boldsymbol{\lambda}}\mathcal{F}(\mathbf{x}, \boldsymbol{\lambda}) - \nabla_{\boldsymbol{\lambda}'}\mathcal{F}(\mathbf{x}', \boldsymbol{\lambda}')\right\|^2 \\
= & \left\|\sum_{i=1}^{m}(\lambda_i - \lambda_i')\mathbf{a}_i\right\|^2 + \left\|\sum_{i=1}^{m}\mathbf{a}_i^\top(\mathbf{x} - \mathbf{x}') + \delta\eta\sum_{i=1}^{m}(\lambda_i' - \lambda_i)\right\|^2 \\
\leq & \|A^\top(\boldsymbol{\lambda} - \boldsymbol{\lambda}')\|^2 + 2\|A(\mathbf{x} - \mathbf{x}')\|^2 + 2\delta^2\eta^2\|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\|^2 \\
\leq & 2\sigma_{\max}^2(A)\|\mathbf{x} - \mathbf{x}'\|^2 + (\sigma_{\max}^2(A) + 2\delta^2\eta^2)\|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\|^2.
\end{aligned}$$

Since

$$\sigma_{\max}(A) = \sqrt{\lambda_{\max}(AA^\top)} \leq \sqrt{\mathrm{Tr}(AA^\top)} \leq \sqrt{m},$$

we have $\sigma_{\max}^2(A) \leq m$, leading to the desired result. ■

Algorithm 2 shows the detailed steps of the mirror prox based algorithm for online convex optimization with long term constraints defined in (5). Compared to Algorithm 1, there are

Algorithm 2 Prox Method with Long Term Constraints

Input: constraints $g_i(\mathbf{x}) \leq 0, i \in [m]$, step size η, and constant δ
Initialization: $\mathbf{z}_1 = \mathbf{0}$ and $\boldsymbol{\mu}_1 = \mathbf{0}$
for $t = 1, 2, \ldots, T$ **do**
  Compute the solution for $\mathbf{x}_t$ and $\boldsymbol{\lambda}_t$ as

$$\mathbf{x}_t = \Pi_{\mathcal{B}} \left(\mathbf{z}_t - \eta \nabla_{\mathbf{x}} \mathcal{F}(\mathbf{z}_t, \boldsymbol{\mu}_t) \right)$$
$$\boldsymbol{\lambda}_t = \Pi_{[0,+\infty)^m} (\boldsymbol{\mu}_t + \eta \nabla_{\boldsymbol{\lambda}} \mathcal{F}(\mathbf{z}_t, \boldsymbol{\mu}_t))$$

  Submit solution $\mathbf{x}_t$
  Receive the convex function $f_t(\mathbf{x})$ and experience loss $f_t(\mathbf{x}_t)$
  Compute $\mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}) = f_t(\mathbf{x}) + \mathcal{F}(\mathbf{x}, \boldsymbol{\lambda}) = f_t(\mathbf{x}) + \sum_{i=1}^m \left\{ \lambda_i g_i(\mathbf{x}) - \frac{\delta \eta}{2} \lambda_i^2 \right\}$
  Update $\mathbf{z}_t$ and $\boldsymbol{\mu}_t$ by

$$\mathbf{z}_{t+1} = \Pi_{\mathcal{B}} \left(\mathbf{z}_t - \eta \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t) \right)$$
$$\boldsymbol{\mu}_{t+1} = \Pi_{[0,+\infty)^m} (\boldsymbol{\mu}_t + \eta \nabla_{\boldsymbol{\lambda}} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t))$$

end for

two key features of Algorithm 2. First, it introduces auxiliary variables $\mathbf{z}_t$ and $\boldsymbol{\mu}_t$ besides the variables $\mathbf{x}_t$ and $\boldsymbol{\lambda}_t$. At each iteration t, it first computes the solutions $\mathbf{x}_t$ and $\boldsymbol{\lambda}_t$ based on the auxiliary variables $\mathbf{z}_t$ and $\boldsymbol{\mu}_t$; it then updates the auxiliary variables based on the gradients computed from $\mathbf{x}_t$ and $\boldsymbol{\lambda}_t$. Second, two different functions are used for updating $(\mathbf{x}_t, \boldsymbol{\lambda}_t)$ and $(\mathbf{z}_t, \boldsymbol{\mu}_t)$: function $\mathcal{F}(\mathbf{x}, \boldsymbol{\lambda})$ is used for computing the solutions $\mathbf{x}_t$ and $\boldsymbol{\lambda}_t$, while function $\mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda})$ is used for updating the auxiliary variables $\mathbf{z}_t$ and $\boldsymbol{\mu}_t$.

Our analysis is based on the Lemma 3.1 from Nemirovski (2005) which is restated here for completeness.

Lemma 10 *Let $\mathsf{B}(\mathbf{x}, \mathbf{x}')$ be a Bregman distance function that has modulus α with respect to a norm $\|\cdot\|$, that is, $\mathsf{B}(\mathbf{x}, \mathbf{x}') \geq \alpha \|\mathbf{x} - \mathbf{x}'\|^2/2$. Given $\mathbf{u} \in \mathcal{B}$, $\mathbf{a}$, and $\mathbf{b}$, we set*

$$\mathbf{w} = \arg\min_{\mathbf{x} \in \mathcal{B}} \mathbf{a}^\top (\mathbf{x} - \mathbf{u}) + \mathsf{B}(\mathbf{x}, \mathbf{u}), \ \mathbf{u}_+ = \arg\min_{\mathbf{x} \in \mathcal{B}} \mathbf{b}^\top (\mathbf{x} - \mathbf{u}) + \mathsf{B}(\mathbf{x}, \mathbf{u}).$$

Then for any $\mathbf{x} \in \mathcal{B}$ and $\eta > 0$, we have

$$\eta \mathbf{b}^\top (\mathbf{w} - \mathbf{x}) \leq \mathsf{B}(\mathbf{x}, \mathbf{u}) - \mathsf{B}(\mathbf{x}, \mathbf{u}_+) + \frac{\eta^2}{2\alpha} \|\mathbf{a} - \mathbf{b}\|_*^2 - \frac{\alpha}{2} \left[\|\mathbf{w} - \mathbf{u}\|^2 + \|\mathbf{w} - \mathbf{u}_+\|^2 \right].$$

We equip $\mathcal{B} \times [0, +\infty)^m$ with the norm $\|\cdot\|$ defined as

$$\|(\mathbf{z}, \boldsymbol{\mu})\|^2 = \frac{\|\mathbf{z}\|^2 + \|\boldsymbol{\mu}\|^2}{2},$$

where $\|\cdot\|^2$ is the Euclidean norm defined separately for each domain. It is immediately seen that the Bregman distance function defined as

$$\mathsf{B}(\mathbf{z}_t, \boldsymbol{\mu}_t, \mathbf{z}_{t+1}, \boldsymbol{\mu}_{t+1}) = \frac{1}{2} \|\mathbf{z}_t - \mathbf{z}_{t+1}\|^2 + \frac{1}{2} \|\boldsymbol{\mu}_t - \boldsymbol{\mu}_{t+1}\|^2$$

is $\alpha = 1$ modules with respect to the norm $\|\cdot\|$.

To analyze the mirror prox algorithm, we begin with a simple lemma which is the direct application of Lemma 10 when applied to the updating rules of Algorithm 3.

Lemma 11 *If $\eta(m + \delta^2\eta^2) \leq \frac{1}{4}$ holds, we have*

$$\begin{aligned}
&\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}) - \mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}_t) \\
&\leq \frac{\|\mathbf{x} - \mathbf{z}_t\|^2 - \|\mathbf{x} - \mathbf{z}_{t+1}\|^2}{2\eta} + \frac{\|\boldsymbol{\lambda} - \boldsymbol{\mu}_t\|^2 - \|\boldsymbol{\lambda} - \boldsymbol{\mu}_{t+1}\|^2}{2\eta} + \eta\|\nabla f_t(\mathbf{x}_t)\|^2.
\end{aligned}$$

Proof To apply Lemma 10, we define $\mathbf{u}$, $\mathbf{w}$, $\mathbf{u}_+$, $\mathbf{a}$ and $\mathbf{b}$ as follows

$$\begin{aligned}
&\mathbf{u} = (\mathbf{z}_t, \boldsymbol{\mu}_t), \mathbf{u}_+ = (\mathbf{z}_{t+1}, \boldsymbol{\mu}_{t+1}), \mathbf{w} = (\mathbf{x}_t, \boldsymbol{\lambda}_t), \\
&\mathbf{a} = (\nabla_{\mathbf{x}}\mathcal{F}(\mathbf{z}_t, \boldsymbol{\mu}_t), -\nabla_{\boldsymbol{\lambda}}\mathcal{F}(\mathbf{z}_t, \boldsymbol{\mu}_t)), \mathbf{b} = (\nabla_{\mathbf{x}}\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t), -\nabla_{\boldsymbol{\lambda}}\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)).
\end{aligned}$$

Using Lemmas 2 and 10, we have

$$\begin{aligned}
&\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}) - \mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}_t) - \frac{\|\mathbf{x} - \mathbf{z}_t\|^2 - \|\mathbf{x} - \mathbf{z}_{t+1}\|^2}{2\eta} - \frac{\|\boldsymbol{\lambda} - \boldsymbol{\mu}_t\|^2 - \|\boldsymbol{\lambda} - \boldsymbol{\mu}_{t+1}\|^2}{2\eta} \\
&\leq \underbrace{\frac{\eta}{2}\left\{\|\nabla_{\mathbf{x}}\mathcal{F}(\mathbf{z}_t, \boldsymbol{\mu}_t) - \nabla_{\mathbf{x}}\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2 + \|\nabla_{\boldsymbol{\lambda}}\mathcal{F}(\mathbf{z}_t, \boldsymbol{\mu}_t) - \nabla_{\boldsymbol{\lambda}}\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2\right\}}_{\mathbf{I}} \\
&- \underbrace{\frac{1}{2}\left\{\|\mathbf{z}_t - \mathbf{x}_t\|^2 + \|\boldsymbol{\mu}_t - \boldsymbol{\lambda}_t\|^2\right\}}_{\mathbf{II}}.
\end{aligned}$$

By expanding the gradient terms and applying the inequality $(a + b)^2 \leq 2(a^2 + b^2)$, we upper bound (**I**) as:

$$\begin{aligned}
(\mathbf{I}) &= \frac{\eta}{2}\{2\|\nabla f_t(\mathbf{x}_t)\|^2 + 2\|\nabla_{\mathbf{x}}\mathcal{F}(\mathbf{z}_t, \boldsymbol{\mu}_t) - \nabla_{\mathbf{x}}\mathcal{F}(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2 + \|\nabla_{\boldsymbol{\lambda}}\mathcal{F}(\mathbf{z}_t, \boldsymbol{\mu}_t) - \nabla_{\boldsymbol{\lambda}}\mathcal{F}(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2\} \\
&\leq \eta\|\nabla f_t(\mathbf{x}_t)\|^2 + \eta\left\{\|\nabla_{\mathbf{x}}\mathcal{F}(\mathbf{z}_t, \boldsymbol{\mu}_t) - \nabla_{\mathbf{x}}\mathcal{F}(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2 + \|\nabla_{\boldsymbol{\lambda}}\mathcal{F}(\mathbf{x}_t, \boldsymbol{\lambda}_t) - \nabla_{\boldsymbol{\lambda}}\mathcal{F}(\mathbf{x}_t, \boldsymbol{\lambda}_t)\|^2\right\} \\
&\leq \eta\|\nabla f_t(\mathbf{x}_t)\|^2 + 2\eta(m + \delta^2\eta^2)\left\{\|\mathbf{z}_t - \mathbf{x}_t\|^2 + \|\boldsymbol{\mu}_t - \boldsymbol{\lambda}_t\|^2\right\}, \qquad (18)
\end{aligned}$$

where the last inequality follows from Proposition 9. Combining (**II**) with (18) results in

$$\begin{aligned}
&\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}) - \mathcal{L}_t(\mathbf{x}, \boldsymbol{\lambda}_t) - \frac{\|\mathbf{x} - \mathbf{z}_t\|^2 - \|\mathbf{x} - \mathbf{z}_{t+1}\|^2}{2\eta} - \frac{\|\boldsymbol{\lambda} - \boldsymbol{\mu}_t\|^2 - \|\boldsymbol{\lambda} - \boldsymbol{\mu}_{t+1}\|^2}{2\eta} \\
&\leq \eta\|\nabla f_t(\mathbf{x}_t)\|_2^2 + \left(2\eta(m + \delta^2\eta^2) - \frac{1}{2}\right)\left\{\|\mathbf{z}_t - \mathbf{x}_t\|^2 + \|\boldsymbol{\mu}_t - \boldsymbol{\lambda}_t\|_2^2\right\}.
\end{aligned}$$

We complete the proof by rearranging the terms and setting $\eta(m + \delta^2\eta^2) \leq \frac{1}{4}$. ■

Theorem 12 *Set $\eta = T^{-1/3}$ and $\delta = T^{-2/3}$. Let $\mathbf{x}_t, t \in [T]$ be the sequence of solutions obtained by Algorithm 2. Then for $T \geq 164(m+1)^3$ we have*

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) \leq O(T^{2/3}) \;\; \textit{and} \;\; \sum_{t=1}^{T} g_i(\mathbf{x}_t) \leq O(T^{2/3}).$$

Proof Similar to the proof of Theorem 4, by summing the bound in Lemma 11 for all rounds $t = 1, \cdots, T$, and taking maximization for $\boldsymbol{\lambda}$ we have the following inequality for any $\mathbf{x}_* \in \mathcal{K}$,

$$\sum_{t=1}^{T} [f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*)] + \sum_{i=1}^{m} \frac{\left[\sum_{t=1}^{T} g_i(\mathbf{x}_t)\right]_+^2}{2(\delta\eta T + m/\eta)} \leq \frac{R^2}{2\eta} + \frac{\eta T}{2} G^2.$$

By setting $\delta = \frac{1}{\eta T}$ and using the fact that $\sum_{t=1}^{T} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) \geq -FT$ we have:

$$\sum_{t=1}^{T} [f_t(\mathbf{x}_t) - f_t(\mathbf{x})] \leq \frac{R^2}{2\eta} + \frac{\eta T}{2} G^2$$

and

$$\sum_{t=1}^{T} g_i(\mathbf{x}_t) \leq \sqrt{(1 + \frac{m}{\eta}) \left(\frac{R^2}{\eta} + \eta T G^2 + FT \right)}.$$

Substituting the stated value for η, we get the desired bounds as mentioned in the theorem. Note that the condition $\eta(m + \delta^2\eta^2) \leq \frac{1}{4}$ in Lemma 11 is satisfied for the stated values of η and δ as long as $T \geq 164(m+1)^3$. ■

Using the same trick as Theorem 8, by introducing appropriate γ, we will be able to establish the solutions that exactly satisfy the constraints in the long run with an $O(T^{2/3})$ regret bound as shown in the following corollary. In the case when all the constraints are linear, that is, $g_i(\mathbf{x}) = \mathbf{a}_i^\top \mathbf{x} \leq b_i, i \in [m]$, Assumption 1 is simplified into the following condition,

$$\min_{\alpha \in \Delta_m} \left\| \sum_{i=1}^{m} \alpha_i \mathbf{a}_i \right\| \geq \sigma, \tag{19}$$

where Δ_m is a m dimensional simplex, that is, $\Delta_m = \{\alpha \in \mathbb{R}_+^m : \sum_{i=1}^m \alpha_i = 1\}$. This is because $g(\mathbf{x}) = \max_{\alpha \in \Delta_m} \sum_{i=1}^m \alpha_i g_i(\mathbf{x})$ and as a result, the (sub)gradient of $g(\mathbf{x})$ can always be written as $\partial g(\mathbf{x}) = \sum_{i=1}^m \alpha_i \nabla g_i(\mathbf{x}) = \sum_{i=1}^m \alpha_i \mathbf{a}_i$ where $\alpha \in \Delta_m$. As an illustrative example, consider the case when the norm vectors $\mathbf{a}_i, i \in [m]$ are linearly independent. In this case the condition mentioned in (19) obviously holds which indicates that the assumption does not limit the applicability of the proposed algorithm.

Corollary 13 *Let $\eta = \delta = T^{-1/3}$. Let $\mathbf{x}_t, t \in [T]$ be the sequence of solutions obtained by Algorithm 2 with $\gamma = bT^{-1/3}$ and $b = 2\sqrt{F}$. With sufficiently large T, that is, $FT \geq R^2 T^{1/3} + G^2 T^{2/3}$, under Assumptions 2 and condition in (19), we have $\mathbf{x}_t, t \in [T]$ satisfy the global constraints $\sum_{t=1}^{T} g_i(\mathbf{x}_t) \leq 0, i \in [m]$ and the regret $\mathfrak{R}_T$ is bounded by*

$$\mathfrak{R}_T = \sum_{t=1}^{T} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) \leq \frac{R^2}{2} T^{1/3} + \left(\frac{G^2}{2} + \frac{2G\sqrt{F}}{\sigma} \right) T^{2/3} = O(T^{2/3}).$$

The proof is similar to that of Theorem 8 and we defer it to Appendix B. As indicated by Corollary 13, for any convex domain defined by a finite number of halfspaces, that is, Polyhedral set, one can easily replace the projection onto the Polyhedral set with the ball containing the Polyhedral at the price of satisfying the constraints in the long run and achieving $O(T^{2/3})$ regret bound.

5. Online Convex Optimization with Long Term Constraints under Bandit Feedback for Domain

We now turn to extending the gradient based convex-concave optimization algorithm discussed in Section 3 to the setting where the learner only receives partial feedback for constraints. More specifically, the exact definition of the domain $\mathcal{K}$ is not exposed to the learner, only that the solution is within a ball $\mathcal{B}$. Instead, after receiving a solution $\mathbf{x}_t$, the oracle will present the learner with the convex loss function $f_t(\mathbf{x})$ and the maximum violation of the constraints for $\mathbf{x}_t$, that is, $g(\mathbf{x}_t) = \max_{i\in[m]} g_i(\mathbf{x}_t)$. We remind that the function $g(\mathbf{x})$ defined in this way is Lipschitz continuous with constant G as proved in Proposition 6. In this setting, the convex-concave function defined in (6) becomes as

$$\mathcal{L}_t(\mathbf{x}, \lambda) = f_t(\mathbf{x}) + \lambda g(\mathbf{x}) - (\delta\eta/2)\lambda^2.$$

The mentioned setting is closely tied to the bandit online convex optimization. In the bandit setting, in contrast to the full information setting, only the cost of the chosen decision (i.e., the incurred loss $f_t(\mathbf{x}_t)$) is revealed to the algorithm, not the function itself. There is a rich body of literature that deals with the bandit online convex optimization. In the seminal papers of Flaxman et al. (2005) and Awerbuch and Kleinberg (2004) it has been shown that one could design algorithms with $O(T^{3/4})$ regret bound even in the bandit setting where only evaluations of the loss functions are revealed at a single point. If we specialize to the online bandit optimization of linear loss functions, Dani et al. (2007) proposed an inefficient algorithm with $O(\sqrt{T})$ regret bound and Abernethy et al. (2008) obtained $O(\sqrt{T \log T})$ bound by an efficient algorithm if the convex set admits an efficiently computable self-concordant barrier. For general convex loss functions, Agarwal et al. (2010) proposed optimal algorithms in a new bandit setting, in which multiple points can be queried for the cost values. By using multiple evaluations, they showed that the modified online gradient descent algorithm can achieve $O(\sqrt{T})$ regret bound in expectation.

Algorithm 3 gives a complete description of the proposed algorithm under the bandit setting, which is a slight modification of Algorithm 1. Algorithm 3 accesses the constraint function $g(\mathbf{x})$ at two points. To facilitate the analysis, we define

$$\widehat{\mathcal{L}}_t(\mathbf{x}, \lambda) = f_t(\mathbf{x}) + \lambda \hat{g}(\mathbf{x}) - \frac{\eta\delta}{2}\lambda^2,$$

where $\hat{g}(\mathbf{x})$ is the smoothed version of $g(\mathbf{x})$ defined as $\hat{g}(\mathbf{x}) = \mathbb{E}_{\mathbf{v}\in\mathbb{S}}[\frac{d}{\zeta} g(\mathbf{x} + \zeta\mathbf{v})\mathbf{v}]$ at point $\mathbf{x}_t$ where $\mathbb{S}$ denotes the unit sphere centered at the origin. Note that $\hat{g}(\mathbf{x})$ is Lipschitz continuous with the same constant G, and it is always differentiable even though $g(\mathbf{x})$ is not in our case.

Since we do not have access to the function $\hat{g}(\cdot)$ to compute $\nabla_{\mathbf{x}}\widehat{\mathcal{L}}(\mathbf{x}, \lambda)$, we need a way to estimate its gradient at point $\mathbf{x}_t$. Our gradient estimation closely follows the idea in

Algorithm 3 Multipoint Bandit Online Convex Optimization with Long Term Constraints

1: **Input**: constraint $g(\mathbf{x})$, step size η, constant $\delta > 0$, exploration parameter $\zeta > 0$, and shrinkage coefficient ξ
2: **Initialization**: $\mathbf{x}_1 = \mathbf{0}$ and $\lambda_1 = 0$
3: **for** $t = 1, 2, \ldots, T$ **do**
4: Submit solution $\mathbf{x}_t$
5: Select unit vector $\mathbf{u}_t$ uniformly at random
6: Query $g(\mathbf{x})$ at points $\mathbf{x}_t + \zeta\mathbf{u}_t$ and $\mathbf{x}_t - \zeta\mathbf{u}_t$ and incur average of them as violation of constraints
7: Compute $\tilde{g}_{\mathbf{x},t} = \nabla f_t(\mathbf{x}_t) + \lambda_t \left[\frac{d}{2\zeta}(g(\mathbf{x}_t + \zeta\mathbf{u}_t) - g(\mathbf{x}_t - \zeta\mathbf{u}_t))\mathbf{u}_t\right]$
8: Compute $\tilde{g}_{\lambda,t} = \frac{1}{2}(g(\mathbf{x}_t + \zeta\mathbf{u}_t) + g(\mathbf{x}_t - \zeta\mathbf{u}_t)) - \eta\delta\lambda_t$
9: Receive the convex function $f_t(\mathbf{x})$ and experience loss $f_t(\mathbf{x}_t)$
10: Update $\mathbf{x}_t$ and λ_t by

$$\mathbf{x}_{t+1} = \Pi_{(1-\xi)\mathcal{B}}\left(\mathbf{x}_t - \eta\tilde{g}_{\mathbf{x},t}\right)$$
$$\lambda_{t+1} = \Pi_{[0,+\infty)}(\lambda_t + \eta\tilde{g}_{\lambda,t})$$

11: **end for**

Agarwal et al. (2010) by querying $g(\mathbf{x})$ function at two points. The main advantage of using two points to estimate the gradient with respect to one point gradient estimation used in Flaxman et al. (2005) is that the former has a bounded norm which is independent of ζ and leads to improved regret bounds.

The gradient estimators for $\nabla_{\mathbf{x}}\widehat{\mathcal{L}}(\mathbf{x}_t, \lambda_t) = \nabla f(\mathbf{x}_t) + \lambda_t \nabla \hat{g}(\mathbf{x}_t)$ and $\nabla_{\lambda}\widehat{\mathcal{L}}(\mathbf{x}_t, \lambda_t) = \hat{g}(\mathbf{x}_t) - \delta\eta\lambda_t$ in Algorithm 3 are computed by evaluating the $g(\mathbf{x})$ function at two random points around $\mathbf{x}_t$ as

$$\tilde{g}_{\mathbf{x},t} = \nabla f_t(\mathbf{x}_t) + \lambda_t \left[\frac{d}{2\zeta}(g(\mathbf{x}_t + \zeta\mathbf{u}_t) - g(\mathbf{x}_t - \zeta\mathbf{u}_t))\mathbf{u}_t\right]$$

and

$$\tilde{g}_{\lambda,t} = \frac{1}{2}(g(\mathbf{x}_t + \zeta\mathbf{u}_t) + g(\mathbf{x}_t - \zeta\mathbf{u}_t)) - \eta\delta\lambda_t,$$

where $\mathbf{u}_t$ is chosen uniformly at random from the surface of the unit sphere. Using Stock's theorem, Flaxman et al. (2005) showed that $\frac{1}{2\zeta}(g(\mathbf{x}_t + \zeta\mathbf{u}_t) - g(\mathbf{x}_t - \zeta\mathbf{u}_t))\mathbf{u}_t$ is a conditionally unbiased estimate of the gradient of $\hat{g}(\mathbf{x})$ at point $\mathbf{x}_t$. To make sure that randomized points around $\mathbf{x}_t$ live inside the convex domain $\mathcal{B}$, we need to stay away from the boundary of the set such that the ball of radius ζ around $\mathbf{x}_t$ is contained in $\mathcal{B}$. In particular, in Flaxman et al. (2005) it has been shown that for any $\mathbf{x} \in (1-\xi)\mathcal{B}$ and any unit vector $\mathbf{u}$ it holds that $(\mathbf{x} + \zeta\mathbf{u}) \in \mathcal{B}$ as soon as $\zeta \in [0, \xi r]$.

In order to facilitate the analysis of the Algorithm 3, we define the convex-concave function $\mathcal{H}_t(\cdot, \cdot)$ as

$$\mathcal{H}_t(\mathbf{x}, \lambda) = \widehat{\mathcal{L}}_t(\mathbf{x}, \lambda) + \left(\tilde{g}_{\mathbf{x},t} - \nabla_{\mathbf{x}}\widehat{\mathcal{L}}(\mathbf{x}_t, \lambda_t)\right)\mathbf{x} + \left(\tilde{g}_{\lambda,t} - \nabla_{\lambda}\widehat{\mathcal{L}}(\mathbf{x}_t, \lambda_t)\right)\lambda. \tag{20}$$

It is easy to check that $\nabla_{\mathbf{x}}\mathcal{H}(\mathbf{x}_t,\lambda_t)=\tilde{g}_{\mathbf{x},t}$ and $\nabla_{\lambda}\mathcal{H}(\mathbf{x}_t,\lambda_t)=\tilde{g}_{\lambda,t}$. By defining functions $\mathcal{H}_t(\mathbf{x},\lambda)$, Algorithm 3 reduces to Algorithm 1 by doing gradient descent on functions $\mathcal{H}_t(\mathbf{x},\lambda)$ except the projection is made onto the set $(1-\xi)\mathcal{B}$ instead of $\mathcal{B}$.

We begin our analysis by reproducing Proposition 3 for functions $\mathcal{H}_t(\cdot,\cdot)$.

Lemma 14 *If the Algorithm 1 is performed over convex set $\mathcal{K}$ with functions $\mathcal{H}_t$ defined in (20), then for any $\mathbf{x}\in\mathcal{K}$ we have*

$$\sum_{t=1}^{T}\mathcal{H}_t(\mathbf{x}_t,\lambda)-\mathcal{H}_t(\mathbf{x},\lambda_t)\leq\frac{R^2+\|\lambda\|_2^2}{2\eta}+\eta(D^2+G^2)T+\eta(d^2G^2+\eta^2\delta^2)\sum_{t=1}^{T}\lambda_t^2.$$

Proof We have $\nabla_{\mathbf{x}}\mathcal{H}_t(\mathbf{x}_t,\lambda_t)=\tilde{g}_{\mathbf{x},t}$ and $\nabla_{\lambda}\mathcal{H}_t(\mathbf{x}_t,\lambda_t)=\tilde{g}_{\lambda,t}$. It is straightforward to show that $\frac{1}{2\zeta}(g(\mathbf{x}_t+\zeta\mathbf{u}_t)-g(\mathbf{x}_t-\zeta\mathbf{u}_t))\mathbf{u}_t$ has norm bounded by Gd (Agarwal et al., 2010). So, the norm of gradients are bounded as $\|\tilde{g}_{\mathbf{x},t}\|_2^2\leq 2(G^2+d^2G^2\lambda_t^2)$ and $\|\tilde{g}_{\lambda,t}\|_2^2\leq 2(D^2+\eta^2\delta^2\lambda_t^2)$. Using Lemma 2, by adding for all rounds we get the desired inequality. ■

The following theorem gives the regret bound and the expected violation of the constraints in the long run for Algorithm 3.

Theorem 15 *Let $c=\sqrt{D^2+G^2}(\sqrt{2}R+\frac{\sqrt{2}D}{\delta R})+(\frac{D}{r}+1)\frac{GD}{r}$. Set $\eta=R/\sqrt{2(D^2+G^2)T}$. Choose δ such that $\delta\geq 2(d^2G^2+\eta^2\delta^2)$. Let $\zeta=\frac{\delta}{T}$ and $\xi=\frac{\zeta}{r}$. Let $\mathbf{x}_t, t\in[T]$ be the sequence of solutions obtained by Algorithm 3. We then have*

$$\sum_{t=1}^{T}f_t(\mathbf{x}_t)-f_t(\mathbf{x})\leq\frac{GD}{r}+c\sqrt{T}=O(T^{1/2}),\ \textit{and}$$
$$\mathbb{E}\Big[\sum_{t=1}^{T}g(\mathbf{x}_t)\Big]\leq G\delta+\sqrt{\Big(\frac{\delta R^2+2(D^2+G^2)}{R\sqrt{D^2+G^2}}\Big)(\frac{GD}{r}+c\sqrt{T}+FT)\sqrt{T}}=O(T^{3/4}).$$

Proof Using Lemma 2 for the functions $\widehat{\mathcal{L}}_t(\cdot,\cdot)$ and $\mathcal{H}_t(\cdot,\cdot)$ we have

$$\widehat{\mathcal{L}}_t(\mathbf{x}_t,\lambda)-\widehat{\mathcal{L}}_t(\mathbf{x},\lambda_t)\leq(\mathbf{x}_t-\mathbf{x})^{\top}\nabla_{\mathbf{x}}\widehat{\mathcal{L}}_t(\mathbf{x}_t,\lambda_t)-(\lambda-\lambda_t)^{\top}\nabla_{\lambda}\widehat{\mathcal{L}}_t(\mathbf{x}_t,\lambda_t),$$

and also

$$\mathcal{H}_t(\mathbf{x}_t,\lambda)-\mathcal{H}_t(\mathbf{x},\lambda_t)\leq(\mathbf{x}_t-\mathbf{x})^{\top}\tilde{g}_{\mathbf{x},t}-(\lambda-\lambda_t)^{\top}\tilde{g}_{\lambda,t}.$$

Subtracting the preceding inequalities, taking expectation, and summing for all t from 1 to T we get

$$\begin{aligned}&\mathbb{E}\left[\sum_{t=1}^{T}\widehat{\mathcal{L}}_t(\mathbf{x}_t,\lambda)-\widehat{\mathcal{L}}_t(\mathbf{x},\lambda_t)\right]\\&=\mathbb{E}\left[\sum_{t=1}^{T}\mathcal{H}_t(\mathbf{x}_t,\lambda)-\mathcal{H}_t(\mathbf{x},\lambda_t)\right]\\&\quad+\mathbb{E}\left[\sum_{t=1}^{T}(\mathbf{x}_t-\mathbf{x})^{\top}(\nabla_{\mathbf{x}}\widehat{\mathcal{L}}_t(\mathbf{x}_t,\lambda_t)-\mathbb{E}_t[\tilde{g}_{\mathbf{x}_t,t}])+(\lambda_t-\lambda)^{\top}(\nabla_{\lambda}\widehat{\mathcal{L}}_t(\mathbf{x}_t,\lambda_t)-\mathbb{E}_t[\tilde{g}_{\lambda_t,t}])\right].\end{aligned}\tag{21}$$

Next we provide an upper bound on the difference between the gradients of two functions. First, $\mathbb{E}_t[\tilde{g}_{\mathbf{x},t}] = \nabla_{\mathbf{x}}\widehat{\mathcal{L}}_t(\mathbf{x}_t, \lambda_t)$, so $\tilde{g}_{\mathbf{x},t}$ is an unbiased estimator of $\nabla_{\mathbf{x}}\widehat{\mathcal{L}}_t(\mathbf{x}_t, \lambda_t)$. Considering the update rule for λ_{t+1} we have $|\lambda_{t+1}| \leq (1-\eta^2\delta)|\lambda_t| + \eta D$ which implies that $|\lambda_t| \leq \frac{D}{\delta\eta}$ for all t. So we obtain

$$\begin{aligned}
&(\lambda_t - \lambda)^\top(\nabla_\lambda \widehat{\mathcal{L}}_t(\mathbf{x}_t, \lambda_t) - \mathbb{E}_t[\tilde{g}_{\lambda_t,t}]) \\
&\leq |\lambda_t - \lambda|\mathbb{E}_t\Big[\|\nabla_\lambda \widehat{\mathcal{L}}_t(\mathbf{x}_t, \lambda_t) - \tilde{g}_{\lambda_t,t}\|_2\Big] \\
&\leq \frac{D}{\delta\eta}\left|\frac{1}{2}(g(\mathbf{x}_t + \zeta\mathbf{u}_t) + g(\mathbf{x}_t - \zeta\mathbf{u}_t)) - \hat{g}(\mathbf{x}_t)\right| \leq \frac{DG}{\delta\eta}\zeta\|\mathbf{u}_t\| \leq \frac{DG}{\delta\eta}\zeta, \qquad (22)
\end{aligned}$$

where the last inequality follows from Lipschitz property of the functions $g(\mathbf{x})$ and $\hat{g}(\mathbf{x})$ with the same constant G. Combining the inequalities (21) and (22) and using Lemma 14, we have

$$\mathbb{E}\Big[\sum_{t=1}^T \widehat{\mathcal{L}}_t(\mathbf{x}_t, \lambda) - \widehat{\mathcal{L}}_t(\mathbf{x}, \lambda_t)\Big] \leq \frac{R^2 + \lambda^2}{2\eta} + \eta(D^2 + G^2)T + \eta(d^2G^2 + \eta^2\delta^2)\sum_{t=1}^T \lambda_t^2 + \frac{DG\zeta}{\delta\eta}T.$$

By expanding the right hand side of above inequality, we obtain

$$\begin{aligned}
&\sum_{t=1}^T [f_t(\mathbf{x}_t) - f_t((1-\xi)\mathbf{x})] + \lambda\mathbb{E}\Big[\sum_{t=1}^T \hat{g}(\mathbf{x}_t)\Big] - \mathbb{E}\Big[\hat{g}((1-\xi)\mathbf{x})\Big]\sum_{t=1}^T \lambda_t - \frac{\eta\delta T}{2}\lambda^2 + \frac{\eta\delta}{2}\sum_{t=1}^T \lambda_t^2 \\
&\leq \frac{R^2 + \lambda^2}{2\eta} + \eta(D^2 + G^2)T + \eta(d^2G^2 + \eta^2\delta^2)\sum_{t=1}^T \lambda_t^2 + \frac{DG\zeta}{\delta\eta}T.
\end{aligned}$$

By choosing $\delta \geq 2(d^2G^2 + \eta^2\delta^2)$ we cancel λ_t^2 terms from both sides and have

$$\begin{aligned}
&\sum_{t=1}^T [f_t(\mathbf{x}_t) - f_t((1-\xi)\mathbf{x})] + \lambda\mathbb{E}\Big[\sum_{t=1}^T \hat{g}(\mathbf{x}_t)\Big] - \mathbb{E}\Big[\hat{g}((1-\xi)\mathbf{x})\Big]\sum_{t=1}^T \lambda_t - \frac{\eta\delta T}{2}\lambda^2 \\
&\leq \frac{R^2 + \lambda^2}{2\eta} + \eta(D^2 + G^2)T + \frac{DG\zeta}{\delta\eta}T. \qquad (23)
\end{aligned}$$

By convexity and Lipschitz property of $f_t(\mathbf{x})$ and $g(\mathbf{x})$ we have

$$f_t((1-\xi)\mathbf{x}) \leq (1-\xi)f_t(\mathbf{x}) + \xi f_t(0) \leq f_t(\mathbf{x}) + DG\xi, \qquad (24)$$

$$g(\mathbf{x}) \leq \hat{g}(\mathbf{x}) + G\zeta \text{ , and } \quad \hat{g}((1-\xi)\mathbf{x}) \leq g((1-\xi)\mathbf{x}) + G\zeta \leq g(\mathbf{x}) + G\zeta + DG\xi. \qquad (25)$$

Plugging (24) and (25) back into (23), for any optimal solution $\mathbf{x}_* \in \mathcal{K}$ we get

$$\begin{aligned}
&\sum_{t=1}^T [f_t(\mathbf{x}_t) - f_t(\mathbf{x})] + \lambda\mathbb{E}\Big[\sum_{t=1}^T g(\mathbf{x}_t)\Big] - \frac{\eta\delta T}{2}\lambda^2 - \lambda G\zeta T \\
&\leq \frac{R^2 + \lambda^2}{2\eta} + \eta(D^2 + G^2)T + \frac{DG\zeta}{\delta\eta}T + DG\xi T + (DG\xi + G\zeta)\sum_{t=1}^T \lambda_t. \qquad (26)
\end{aligned}$$

Considering the fact that $\lambda_t \leq \frac{D}{\delta\eta}$ we have $\sum_{t=1}^{T} \lambda_t \leq \frac{DT}{\delta\eta}$. Plugging back into the (26) and rearranging the terms we have

$$\sum_{t=1}^{T} [f_t(\mathbf{x}_t) - f_t(\mathbf{x})] + \lambda \mathbb{E}\Big[\sum_{t=1}^{T} g(\mathbf{x}_t)\Big] - \frac{\eta\delta T}{2}\lambda^2 - \lambda G\zeta T - \frac{\lambda^2}{2\eta}$$
$$\leq \frac{R^2}{2\eta} + \eta(D^2 + G^2)T + \frac{DG\zeta}{\delta\eta}T + DG\xi T + (DG\xi + G\zeta)\frac{DT}{\delta\eta}.$$

By setting $\xi = \frac{\zeta}{r}$ and $\zeta = \frac{1}{T}$ we get

$$\sum_{t=1}^{T} [f_t(\mathbf{x}_t) - f_t(\mathbf{x})] \leq \frac{R^2}{2\eta} + \eta(D^2 + G^2)T + \frac{DG\zeta}{\delta\eta}T + \frac{\zeta DGT}{r} + (\frac{D}{r} + 1)\zeta\frac{DGT}{\delta\eta},$$

which gives the mentioned regret bound by optimizing for η. Maximizing for λ over the range $(0, +\infty)$ and using $\sum_{t=1}^{T} f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) \geq -FT$, yields the following inequality for the violation of constraints

$$\frac{\Big[\mathbb{E}\left[\sum_{t=1}^{T} g(\mathbf{x}_t)\right] - G\zeta T\Big]_+^2}{4(\delta\eta T/2 + 1/2\eta)} \leq \frac{DG}{r} + c\sqrt{T} + FT.$$

Plugging in the stated values of parameters completes the proof. Note that $\delta = 4d^2G^2$ obeys the condition specified in the theorem. ■

6. Conclusion

In this study we have addressed the problem of online convex optimization with constraints, where we only need the constraints to be satisfied in the long run. In addition to the regret bound which is the main tool in analyzing the performance of general online convex optimization algorithms, we defined the bound on the violation of constraints in the long term which measures the cumulative violation of the solutions from the constraints for all rounds. Our setting is applied to solving online convex optimization without projecting the solutions onto the complex convex domain at each iteration, which may be computationally inefficient for complex domains. Our strategy is to turn the problem into an online convex-concave optimization problem and apply online gradient descent algorithm to solve it. We have proposed efficient algorithms in three different settings; the violation of constraints is allowed, the constraints need to be exactly satisfied, and finally we do not have access to the target convex domain except it is bounded by a ball. Moreover, for domains determined by linear constraints, we used the mirror prox method, a simple gradient based algorithm for variational inequalities, and obtained an $O(T^{2/3})$ bound for both regret and the violation of the constraints.

Our work leaves open a number of interesting directions for future work. In particular it would be interesting to see if it is possible to improve the bounds obtained in this paper, i.e., getting an $O(\sqrt{T})$ bound on the regret and better bound than $O(T^{3/4})$ on the violation

of constraints for general convex domains. Proving optimal lower bounds for the proposed setting also remains as an open question. Also, it would be interesting to consider strongly convex loss or constraint functions. Finally, relaxing the assumption we made to exactly satisfy the constraints in the long run is an interesting problem to be investigated.

Acknowledgments

The authors would like to thank the Action Editor and three anonymous reviewers for their constructive comments and helpful suggestions on the original version of this paper. This work was supported in part by National Science Foundation (IIS-0643494) and Office of Navy Research (Award N000141210431).

Appendix A. Proof of Theorem 1

We first show that when $\delta < 1$, there exists a loss function and a constraint function such that the violation of constraint is linear in T. To see this, we set $f_t(\mathbf{x}) = \mathbf{w}^\top \mathbf{x}, t \in [T]$ and $g(\mathbf{x}) = 1 - \mathbf{w}^\top \mathbf{x}$. Assume we start with an infeasible solution, that is, $g(\mathbf{x}_1) > 0$ or $\mathbf{x}_1^\top \mathbf{w} < 1$. Given the solution $\mathbf{x}_t$ obtained at tth trial, using the standard gradient descent approach, we have $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta(1-\delta)\mathbf{w}$. Hence, if $\mathbf{x}_t^\top \mathbf{w} < 1$, since we have $\mathbf{x}_{t+1}^\top \mathbf{w} < \mathbf{x}_t^\top \mathbf{w} < 1$, if we start with an infeasible solution, all the solutions obtained over the trails will violate the constraint $g(\mathbf{x}) \leq 0$, leading to a linear number of violation of constraints. Based on this analysis, we assume $\delta > 1$ in the analysis below.

Given a strongly convex loss function $f(\mathbf{x})$ with modulus γ, we consider a constrained optimization problem given by

$$\min_{g(\mathbf{x}) \leq 0} f(\mathbf{x}),$$

which is equivalent to the following unconstrained optimization problem

$$\min_{\mathbf{x}} f(\mathbf{x}) + \lambda [g(\mathbf{x})]_+,$$

where $\lambda \geq 0$ is the Lagrangian multiplier. Since we can always scale $f(\mathbf{x})$ to make $\lambda \leq 1/2$, it is safe to assume $\lambda \leq 1/2 < \delta$. Let $\mathbf{x}_*$ and $\mathbf{x}_a$ be the optimal solutions to the constrained optimization problems $\arg\min_{g(\mathbf{x}) \leq 0} f(\mathbf{x})$ and $\arg\min_{\mathbf{x}} f(\mathbf{x}) + \delta [g(\mathbf{x})]_+$, respectively. We choose $f(\mathbf{x})$ such that $\|\nabla f(\mathbf{x}_*)\| > 0$, which leads to $\mathbf{x}_a \neq \mathbf{x}_*$. This holds because according to the first order optimality condition, we have

$$\nabla f(\mathbf{x}_*) = -\lambda \nabla g(\mathbf{x}_*), \ \nabla f(\mathbf{x}_a) = -\delta \nabla g(\mathbf{x}_*),$$

and therefore $\nabla f(\mathbf{x}_*) \neq \nabla f(\mathbf{x}_a)$ when $\lambda < \delta$. Define $\Delta = f(\mathbf{x}_a) - f(\mathbf{x}_*)$. Since $\Delta \geq \gamma \|\mathbf{x}_a - \mathbf{x}_*\|^2/2$ due to the strong convexity of $f(\mathbf{x})$, we have $\Delta > 0$.

Let $\{\mathbf{x}_t\}_{t=1}^T$ be the sequence of solutions generated by the OGD algorithm that minimizes the modified loss function $f(\mathbf{x}) + \delta[g(\mathbf{x})]_+$. We have

$$\begin{aligned}
& \sum_{t=1}^{T} f(\mathbf{x}_t) + \delta[g(\mathbf{x}_t)]_+ \geq T \min_{\mathbf{x}} f(\mathbf{x}) + \delta[g(\mathbf{x})]_+ \\
= \; & T(f(\mathbf{x}_a) + \delta[g(\mathbf{x}_a)]_+) \geq T(f(\mathbf{x}_a) + \lambda[g(\mathbf{x}_a)]_+) \\
= \; & T(f(\mathbf{x}_*) + \lambda[g(\mathbf{x}_*)]_+) + T(f(\mathbf{x}_a) + \lambda[g(\mathbf{x}_a)]_+ - f(\mathbf{x}_*) - \lambda[g(\mathbf{x}_*)]) \\
\geq \; & T \min_{g(\mathbf{x}) \leq 0} f(\mathbf{x}) + T\Delta.
\end{aligned}$$

As a result, we have

$$\sum_{t=1}^{T} f(\mathbf{x}_t) + \delta[g(\mathbf{x}_t)]_+ - \min_{g(\mathbf{x}) \leq 0} f(\mathbf{x}) = O(T),$$

implying that either the regret $\sum_{t=1}^T f(\mathbf{x}_t) - T f(\mathbf{x}_*)$ or the violation of the constraints $\sum_{t=1}^T [g(\mathbf{x})]_+$ is linear in T.

To better understand the performance of penalty based approach, here we analyze the performance of the OGD in solving the online optimization problem in (3). The algorithm is analyzed using the following lemma from Zinkevich (2003).

Lemma 16 *Let $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_T$ be the sequence of solutions obtained by applying OGD on the sequence of bounded convex functions $f_1, f_2, \ldots, f_T$. Then, for any solution $\mathbf{x}_* \in \mathcal{K}$ we have*

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{x}_*) \leq \frac{R^2}{2\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \|\nabla f_t(\mathbf{x}_t)\|^2.$$

We apply OGD to functions $\hat{f}_t(\mathbf{x}),\ t \in [T]$ defined in (4), that is, instead of updating the solution based on the gradient of $f_t(\mathbf{x})$, we update the solution by the gradient of $\hat{f}_t(\mathbf{x})$. Using Lemma 16, by expanding the functions $\hat{f}_t(\mathbf{x})$ based on (4) and considering the fact that $\sum_{i=1}^m [g_i(\mathbf{x}_*)]_+^2 = 0$, we get

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{x}_*) + \frac{\delta}{2} \sum_{t=1}^{T} \sum_{i=1}^{m} [g_i(\mathbf{x})]_+^2 \leq \frac{R^2}{2\eta} + \frac{\eta}{2} \sum_{t=1}^{T} \|\nabla \hat{f}_t(\mathbf{x}_t)\|^2. \tag{27}$$

From the definition of $\hat{f}_t(\mathbf{x})$, the norm of the gradient $\nabla \hat{f}_t(\mathbf{x}_t)$ is bounded as follows

$$\|\nabla \hat{f}_t(\mathbf{x})\|^2 = \|\nabla f_t(\mathbf{x}) + \delta \sum_{i=1}^{m} [g_i(\mathbf{x})]_+ \nabla g_i(\mathbf{x})\|^2 \leq 2G^2(1 + m\delta^2 D^2), \tag{28}$$

where the inequality holds because $(a_1 + a_2)^2 \leq 2(a_1^2 + a_2^2)$. By substituting (28) into the (27) we have:

$$\sum_{t=1}^{T} f_t(\mathbf{x}_t) - \sum_{t=1}^{T} f_t(\mathbf{x}_*) + \frac{\delta}{2} \sum_{t=1}^{T} \sum_{i=1}^{m} [g_i(\mathbf{x}_t)]_+^2 \leq \frac{R^2}{2\eta} + \eta G^2(1 + m\delta^2 D^2)T. \tag{29}$$

Since $[\cdot]_+^2$ is a convex function, from Jensen's inequality and following the fact that $\sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_*) \geq -FT$, we have:

$$\frac{\delta}{2T}\sum_{i=1}^m\left[\sum_{t=1}^T g_i(\mathbf{x}_t)\right]_+^2 \leq \frac{\delta}{2}\sum_{i=1}^m\sum_{t=1}^T [g_i(\mathbf{x}_t)]_+^2 \leq \frac{R^2}{2\eta} + \eta G^2(1+m\delta^2D^2)T + FT.$$

By minimizing the right hand side of (29) with respect to η, we get the regret bound as

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - \sum_{t=1}^T f_t(\mathbf{x}_*) \leq RG\sqrt{2(1+m\delta^2D^2)T} = O(\delta\sqrt{T}) \tag{30}$$

and the bound for the violation of constraints as

$$\sum_{t=1}^T g_i(\mathbf{x}_t) \leq \sqrt{\left(\frac{R^2}{2\eta} + \eta G^2(1+m\delta^2D^2)T + FT\right)\frac{2T}{\delta}} = O(T^{1/4}\delta^{1/2} + T\delta^{-1/2}). \tag{31}$$

Examining the bounds obtained in (30) and (31), it turns out that in order to recover $O(\sqrt{T})$ regret bound, we need to set δ to be a constant, leading to $O(T)$ bound for the violation of constraints in the long run, which is not satisfactory at all.

Appendix B. Proof of Corollary 13

Let $\mathbf{x}_\gamma$ be the optimal solution to $\min_{g(\mathbf{x})\leq-\gamma}\sum_{t=1}^T f_t(\mathbf{x})$. Similar to the proof of Theorem 12, we have

$$\sum_{t=1}^T [f_t(\mathbf{x}_t) - f_t(\mathbf{x}_\gamma)] + \frac{\left[\sum_{t=1}^T g(\mathbf{x}_t) + \gamma T\right]_+^2}{2(\delta\eta T + 1/\eta)} \leq \frac{R^2}{2\eta} + \frac{\eta T}{2}G^2.$$

Using the stated values for the parameters $\eta = \delta = T^{-1/3}$, and applying the fact that $\sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_\gamma) \geq -FT$ we obtain,

$$\sum_{t=1}^T f_t(\mathbf{x}_t) - f_t(\mathbf{x}_\gamma) \leq \frac{R^2}{2}T^{1/3} + \frac{G^2}{2}T^{2/3} \tag{32}$$

and

$$\left[\sum_{t=1}^T g(\mathbf{x}_t) + \gamma T\right]_+^2 \leq 2\big(R^2T^{1/3} + G^2T^{2/3} + FT\big)T^{1/3}. \tag{33}$$

From Theorem 7, we have the bound

$$\sum_{t=1}^T f_t(\mathbf{x}_\gamma) \leq \sum_{t=1}^T f_t(\mathbf{x}_*) + \frac{G}{\sigma}\gamma T. \tag{34}$$

Combining inequalities (32) and (34) with substituting the stated value of $\gamma = bT^{-1/3}$ yields the regret bound as desired. To obtain the bound for the violation of the constraints, from (33) we have

$$\sum_{t=1}^{T} g(\mathbf{x}_t) \leq \sqrt{2\big(R^2T^{1/3} + G^2T^{2/3} + FT\big)T^{1/3}} - bT^{2/3}.$$

For sufficiently large values of T, that is, $FT \geq R^2T^{1/3} + G^2T^{2/3}$ we can simplify above inequality as $\sum_{t=1}^{T} g(\mathbf{x}_t) \leq 2\sqrt{F}T^{2/3} - bT^{2/3}$. By setting $b = 2\sqrt{F}$ the zero bound on the violation of constraints is guaranteed.

References

Jacob Abernethy, Elad Hazan, and Alexander Rakhlin. Competing in the dark: An efficient algorithm for bandit linear optimization. In *COLT*, pages 263–274, 2008.

Jacob Abernethy, Alekh Agarwal, Peter L. Bartlett, and Alexander Rakhlin. A stochastic view of optimal regret through minimax duality. In *COLT*, 2009.

Jacob Abernethy, Elad Hazan, and Alexander Rakhlin. Interior-point methods for full-information and bandit online learning. *IEEE Transactions on Information Theory*, 58 (7):4164–4175, 2012.

Alekh Agarwal, Ofer Dekel, and Lin Xiao. Optimal algorithms for online convex optimization with multi-point bandit feedback. In *COLT*, pages 28–40, 2010.

Baruch Awerbuch and Robert D. Kleinberg. Adaptive routing with end-to-end feedback: distributed learning and geometric approaches. In *STOC*, pages 45–53, 2004.

Peter L. Bartlett, Elad Hazan, and Alexander Rakhlin. Adaptive online gradient descent. In *NIPS*, pages 257–269, 2007.

Andrey Bernstein, Shie Mannor, and Nahum Shimkin. Online classification with specificity constraints. In *NIPS*, pages 190–198, 2010.

Dimitri P. Bertsekas, Angelia Nedic, and Asuman E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, 2003.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

Nicolò Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.

Varsha Dani, Thomas P. Hayes, and Sham Kakade. The price of bandit information for online optimization. In *NIPS*, 2007.

John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the l_1-ball for learning in high dimensions. In *ICML*, pages 272–279, 2008.

Abraham Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. In *SODA*, pages 385–394, 2005.

Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

Jun Liu and Jieping Ye. Efficient euclidean projections in linear time. In *ICML*, pages 83–90, 2009.

Shie Mannor and John N. Tsitsiklis. Online learning with constraints. In *COLT*, pages 529–543, 2006.

Shie Mannor, John N. Tsitsiklis, and Jia Yuan Yu. Online learning with sample path constraints. *Journal of Machine Learning Research*, 10:569–590, 2009.

Hariharan Narayanan and Alexander Rakhlin. Random walk approach to regret minimization. In *NIPS*, pages 1777–1785, 2010.

Arkadi Nemirovski. Efficient methods in convex programming. Lecture Notes, Available at http://www2.isye.gatech.edu/ nemirovs, 1994.

Arkadi Nemirovski. Prox-method with rate of convergence o(1/t) for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM J. on Optimization*, 15(1):229–251, 2005.

Alexander Rakhlin. Lecture notes on online learning. Lecture Notes, Available at http://www-stat.wharton.upenn.edu/ rakhlin/papers, 2009.

Shai Shalev-Shwartz and Sham M. Kakade. Mind the duality gap: Logarithmic regret algorithms for online optimization. In *NIPS*, pages 1457–1464, 2008.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, pages 928–936, 2003.